MATLAB EXPO 2019

AI Techniques in MATLAB for Signal, Time-Series, and Text Data

Sylvain Lacaze

AI and Deep Learning for Signals in the News

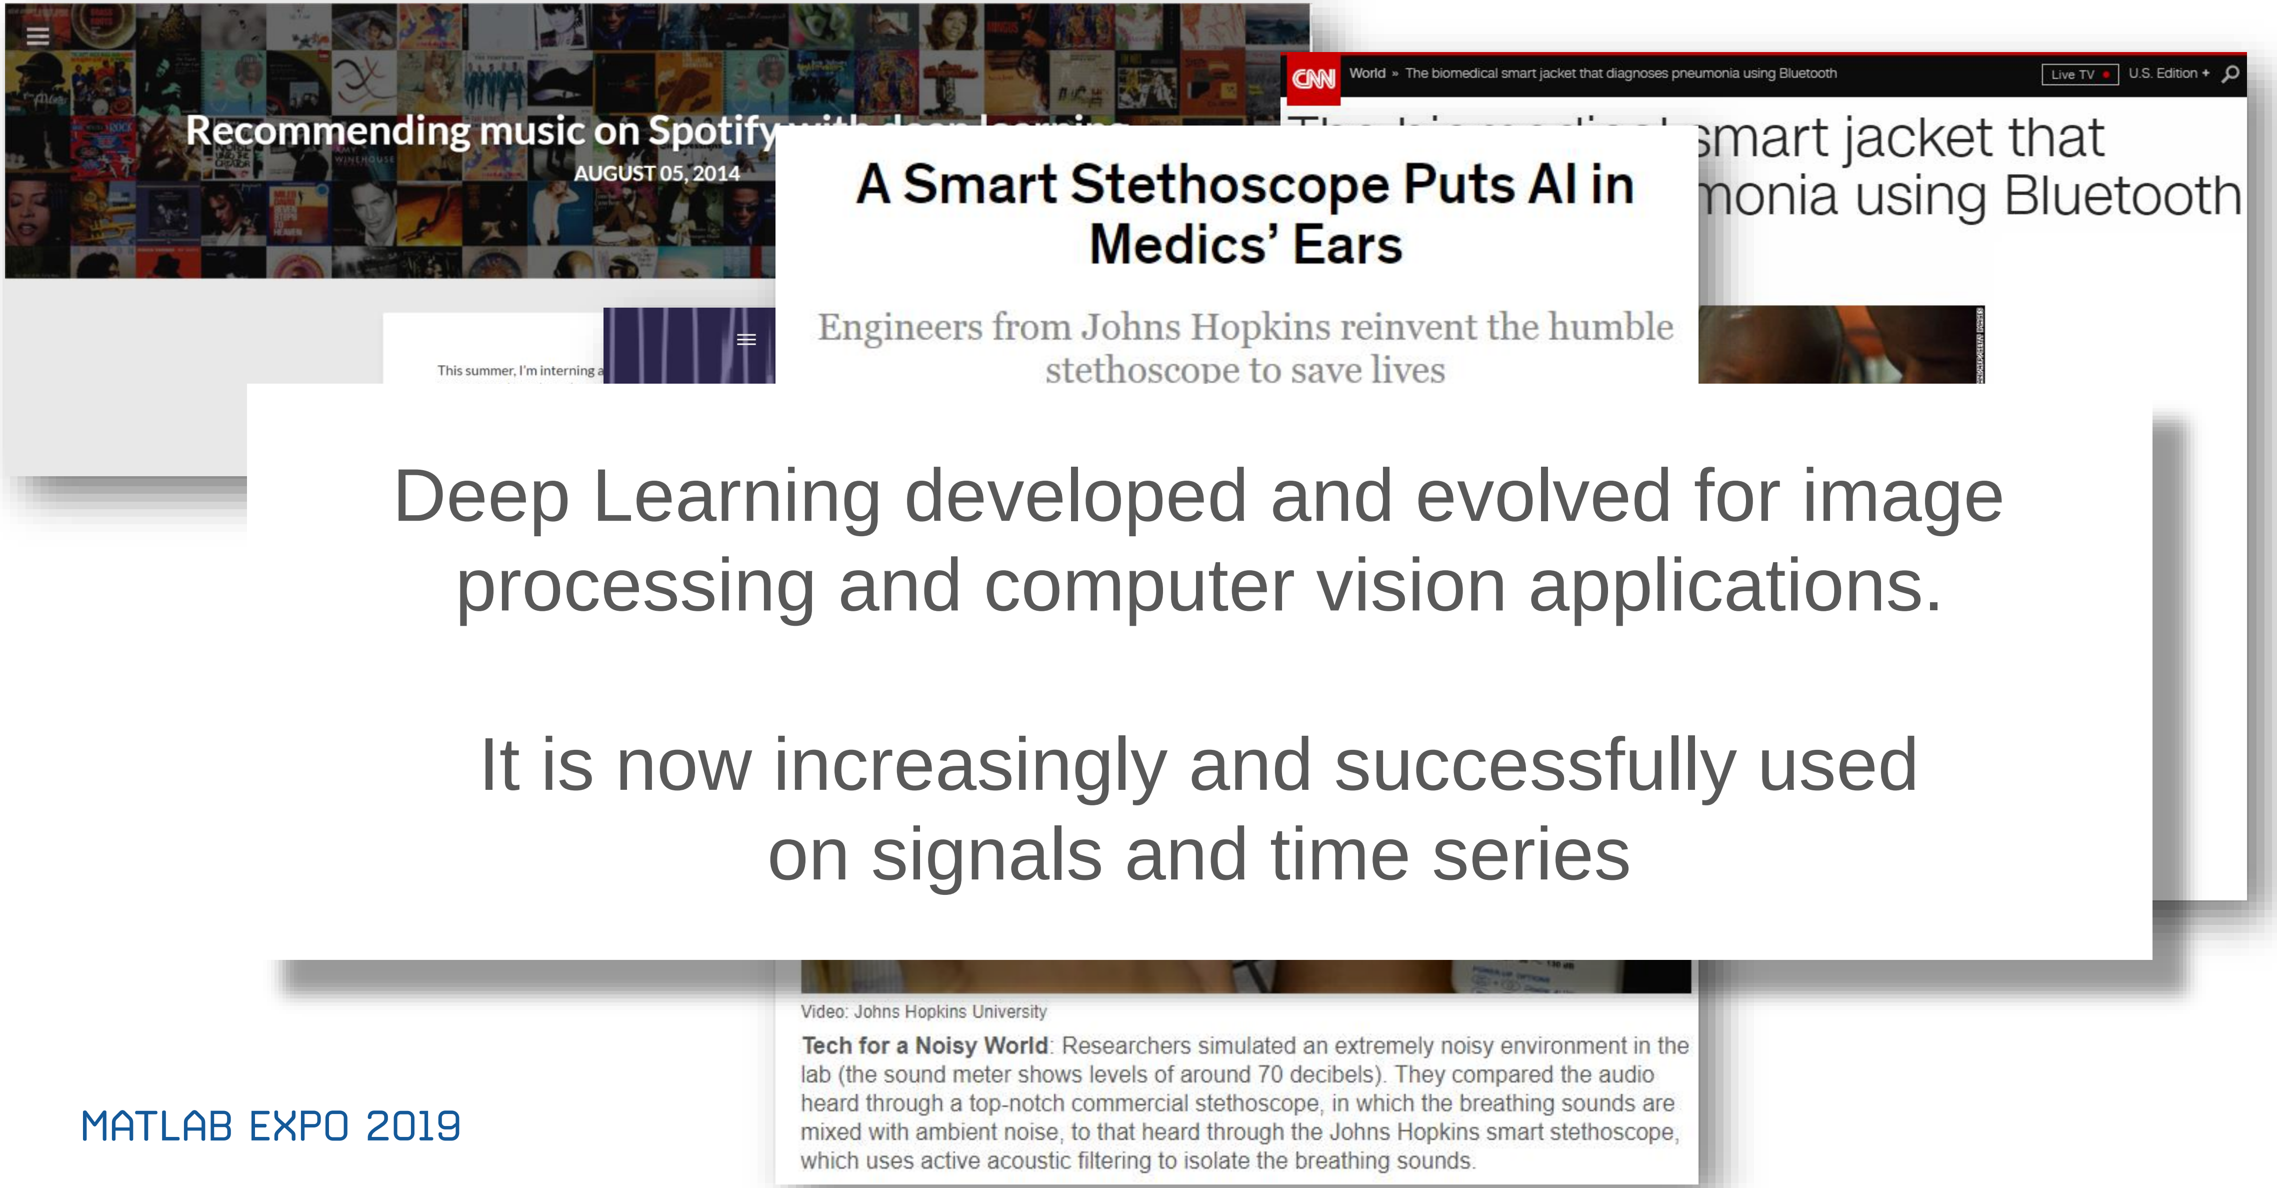

Deep Learning developed and evolved for image processing and computer vision applications.

It is now increasingly and successfully used on signals and time series

The Use of Deep Learning is Growing Across Industries

Aerospace, Defense and Communications

Communications devices, security

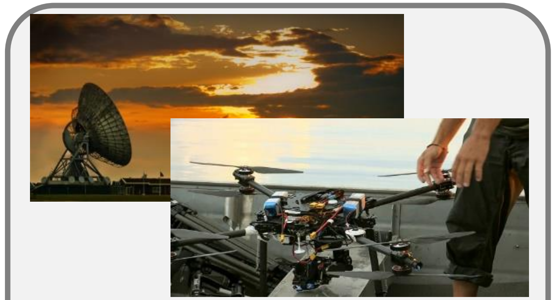

Multi-standard communications receivers, drone recognition

Consumer Electronics and Digital Health

Voice assistants

Digital health

Automotive

Voice control enabled Infotainment

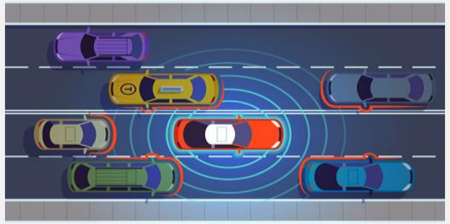

Sensor processing, automated driving

Industrial Automation

Condition monitoring

Predictive maintenance

Application Examples Using MATLAB – Audio and Speech

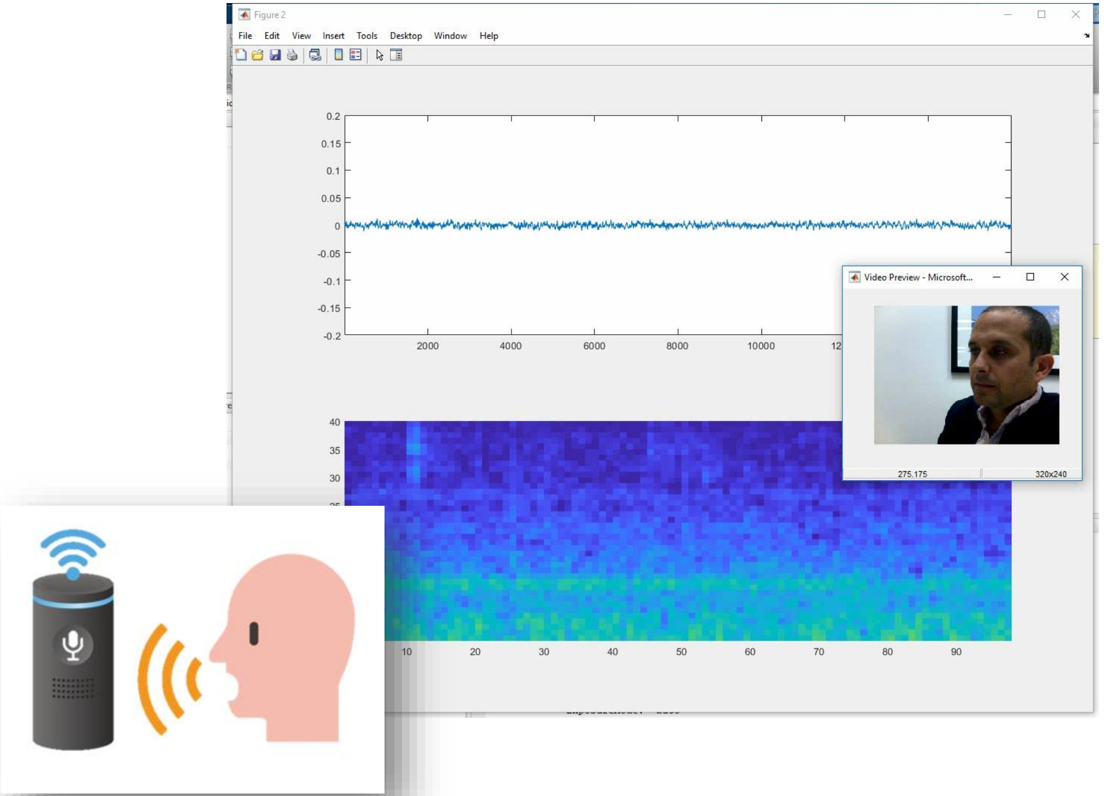

Speech Command Recognition
(a.k.a. "Keyword Spotting")

https://www.mathworks.com/help/deeplearning/examples/deep-learning-speech-recognition.html

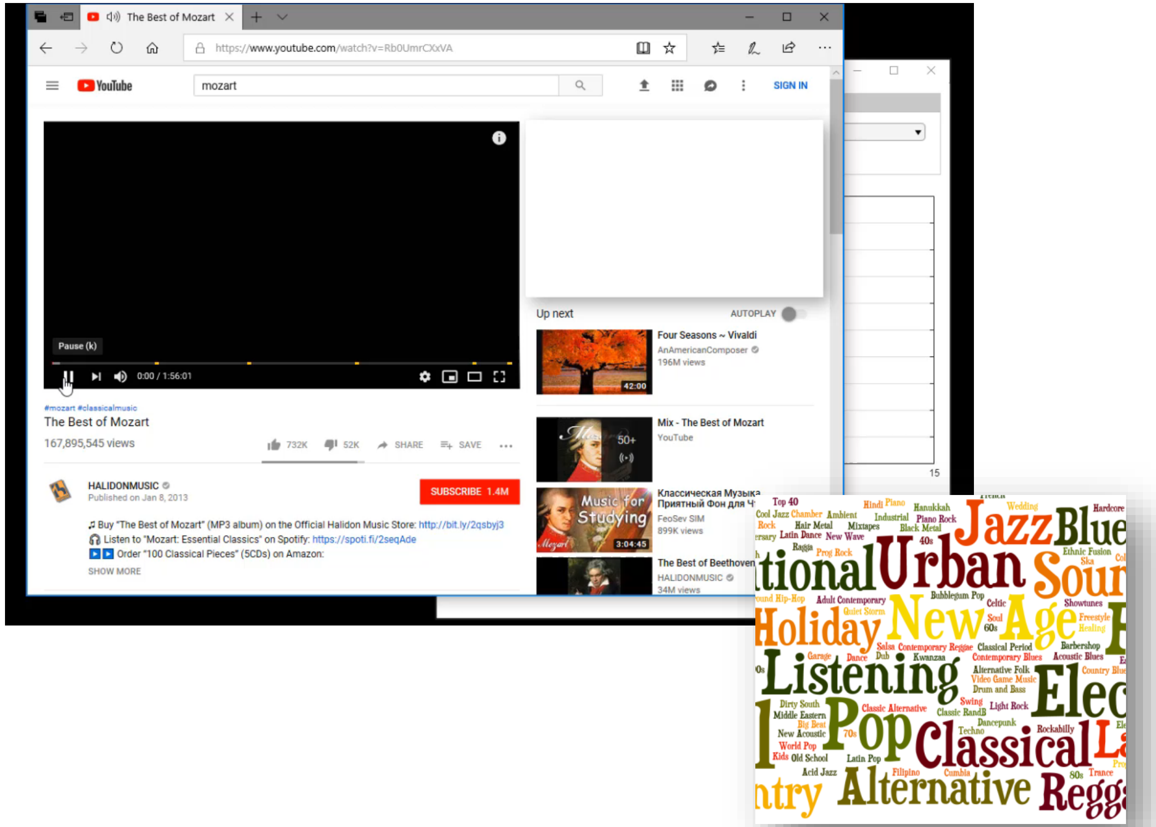

Music Genre Classification

https://www.mathworks.com/help/audio/examples/music-genre-classification-using-wavelet-time-scattering.html

Application Examples Using MATLAB – Industrial and physiological sensors

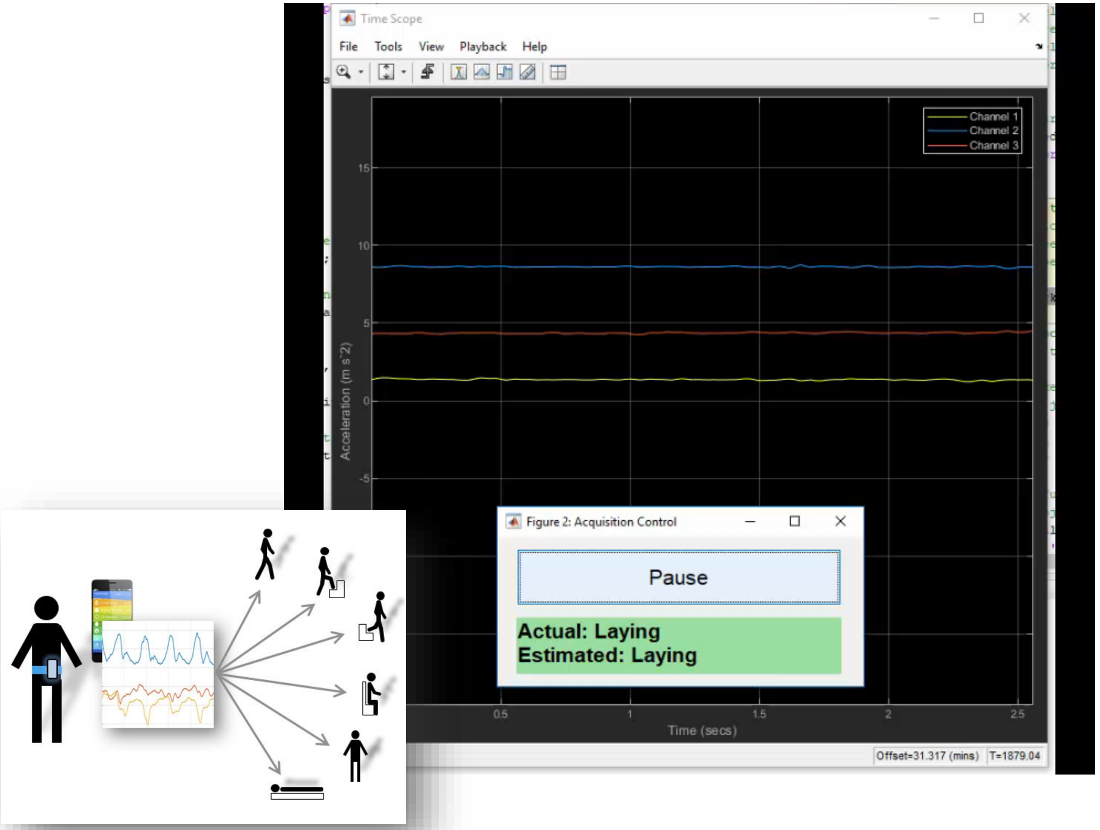

Human Activity Recognition

https://www.mathworks.com/help/deeplearning/examples/sequence-to-sequence-classification-using-deep-learning.html

ECG Signal Classification

https://www.mathworks.com/help/signal/examples/classify-ecg-signals-using-long-short-term-memory-networks.html

Application Examples Using MATLAB – Radar and Communications

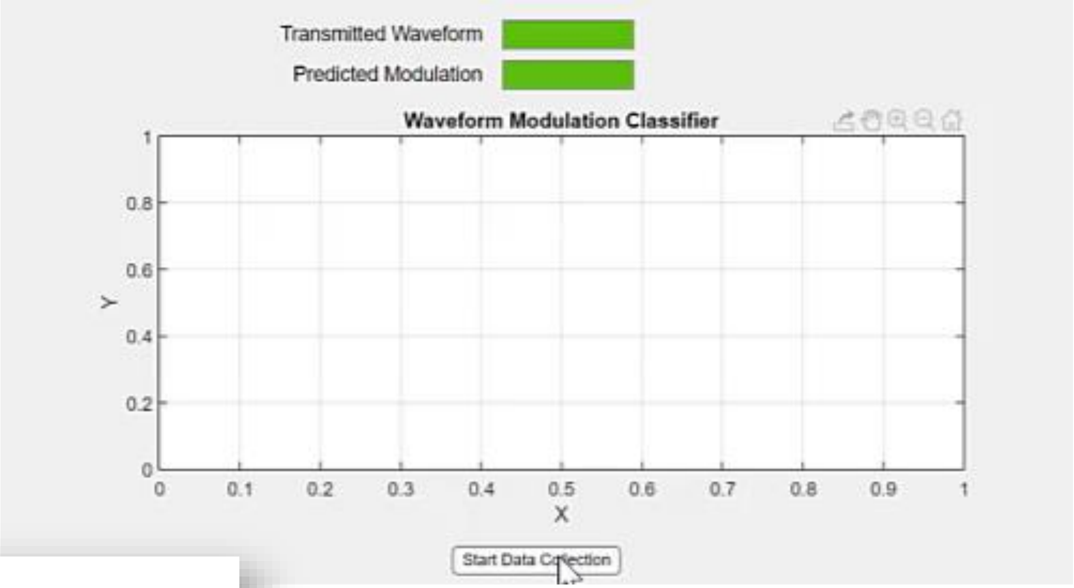

Modulation Classification

http://www.mathworks.com/help/comm/examples/modulation-classification-with-deep-learning.html

Confusion Matrix for Test Data

True Class \ Predicted Class	16QAM	64QAM	8PSK	B-FM	BPSK	CPFSK	GFSK	PAM4	QPSK		
16QAM	89	11								89.0%	11.0%
64QAM	1	99								99.0%	1.0%
8PSK			100							100.0%	
B-FM				100						100.0%	
BPSK					100					100.0%	
CPFSK						100				100.0%	
GFSK							100			100.0%	
PAM4								100		100.0%	
QPSK			4						96	96.0%	4.0%

Agenda

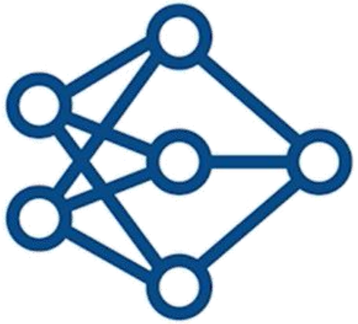

- Deep Learning – Basic ideas
- Deep Learning Model Development for Signals, Time Series, and Text
- Conclusions

What is Deep Learning?

Deep learning is a type of machine learning in which a model learns from examples.

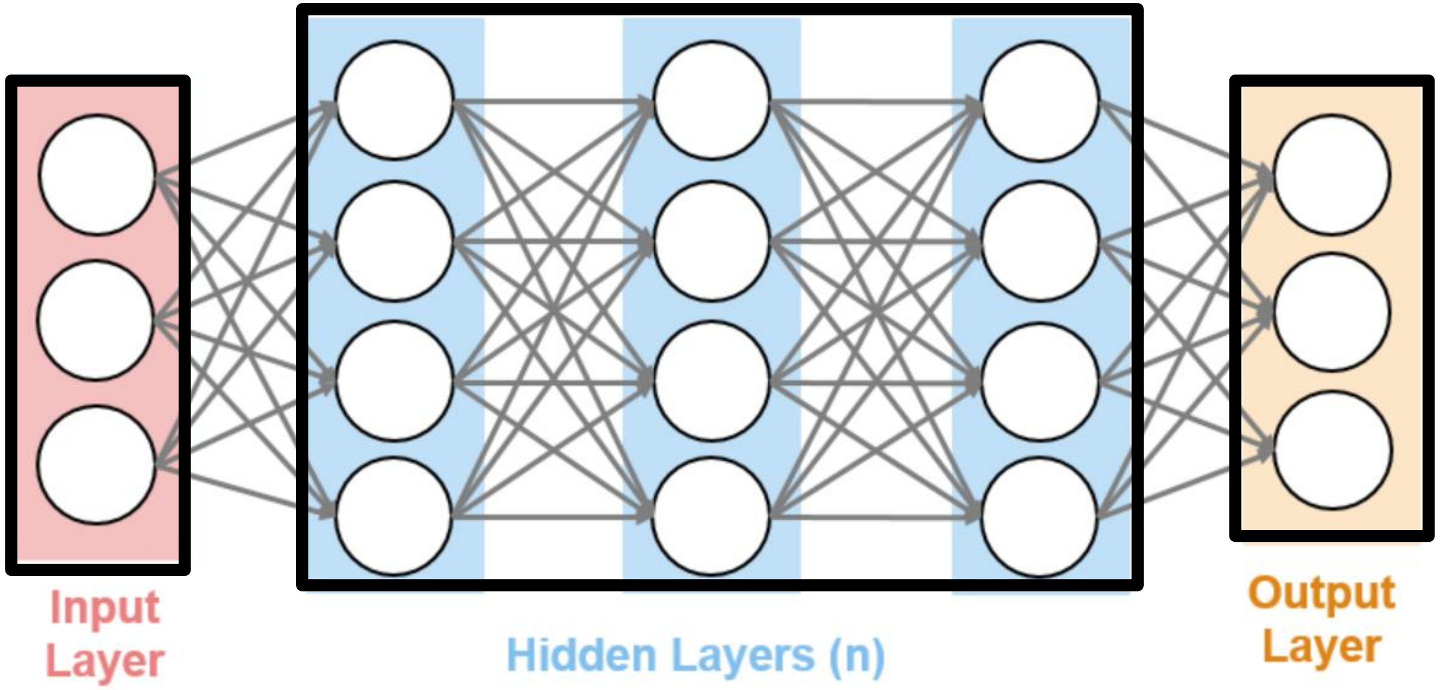

Common Network Architectures - Signal Processing

Convolutional Neural Networks (**CNN**)

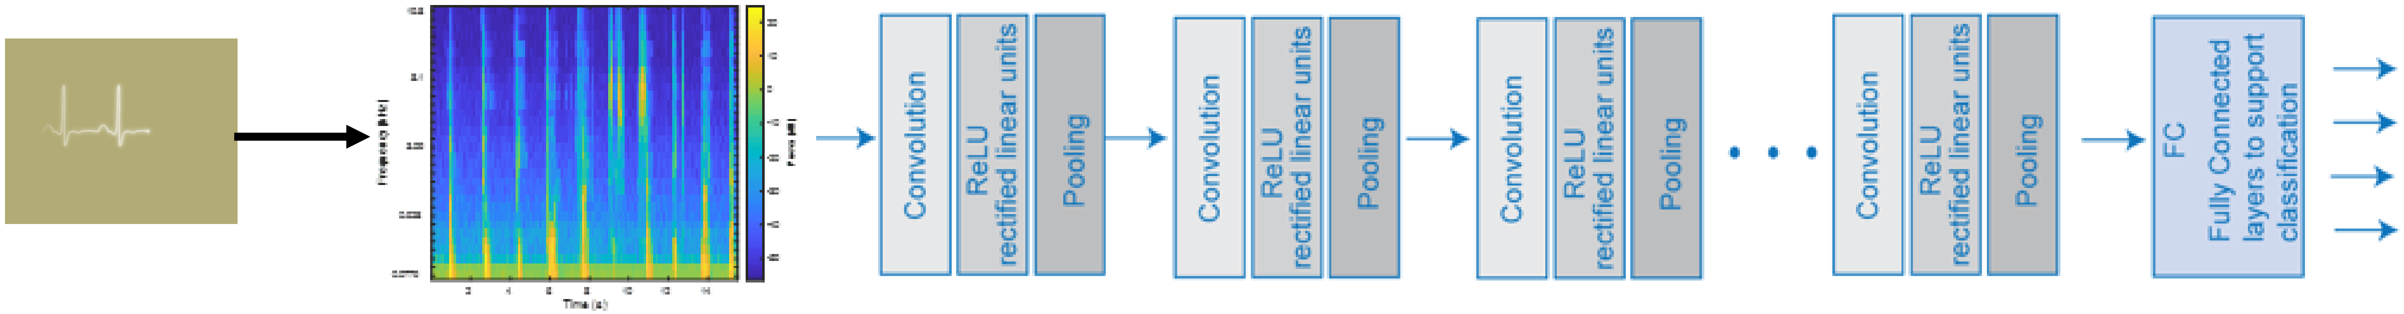

Time-Frequency Transformation

Long Short Term Memory (**LSTM**) Networks

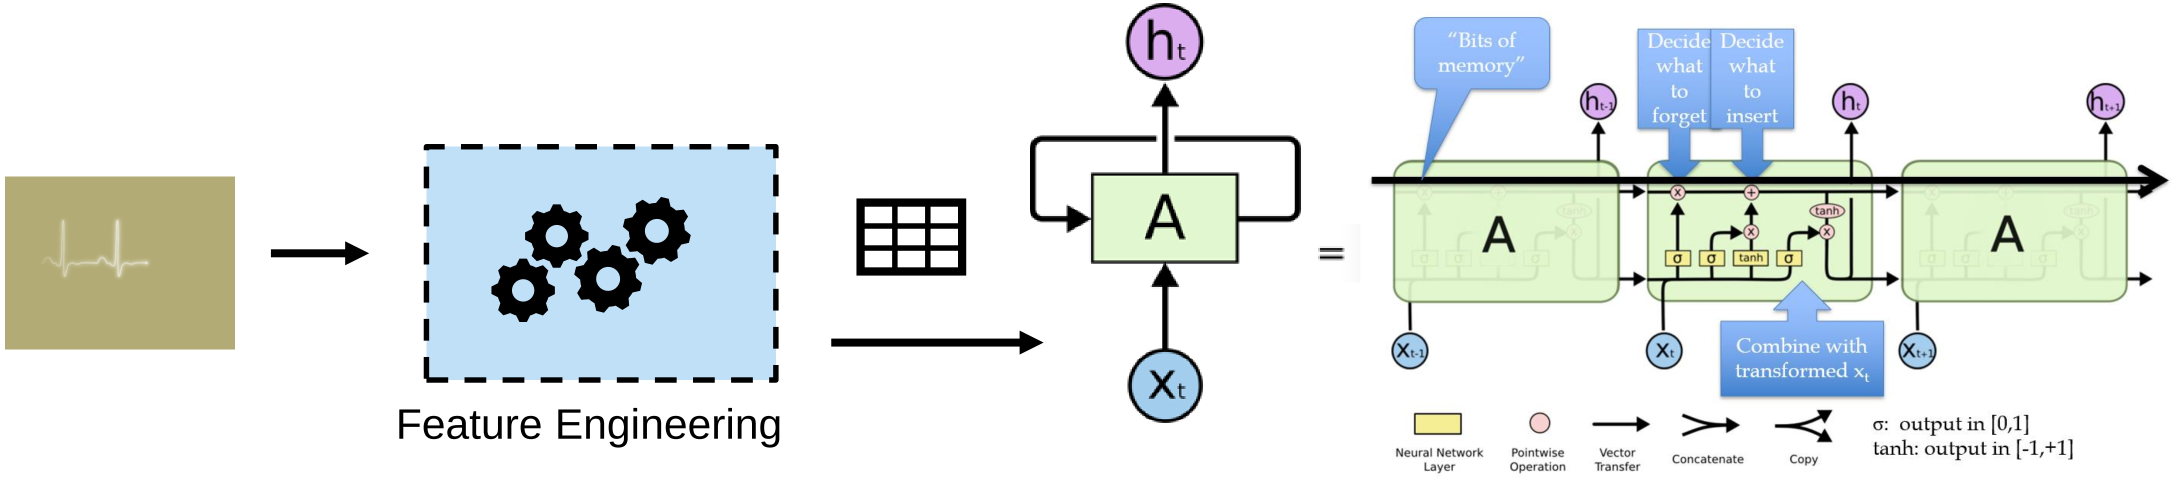

Feature Engineering

Common Network Architectures – Text Analytics

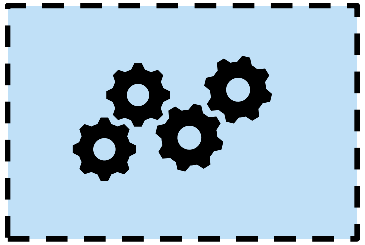

Feature Engineering

Convolutional Neural Networks (**CNN**)

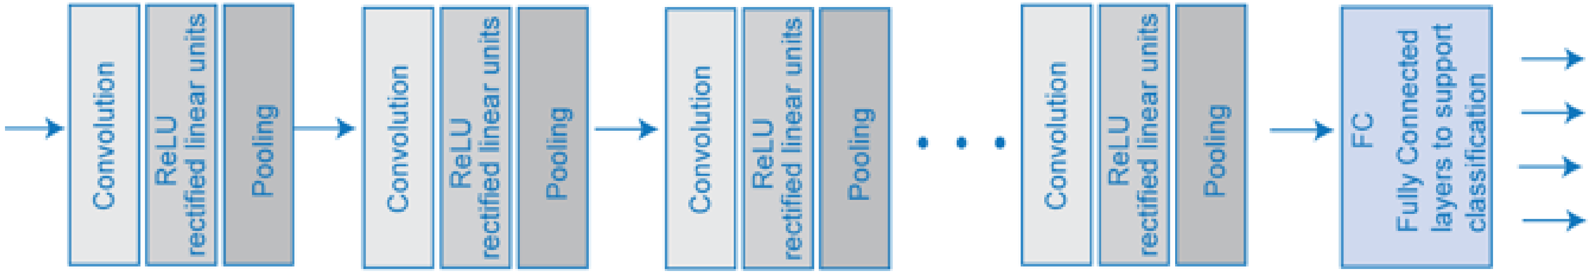

Long Short Term Memory (**LSTM**) Networks

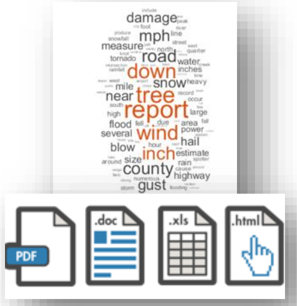

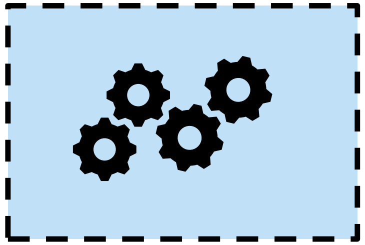

Feature Engineering

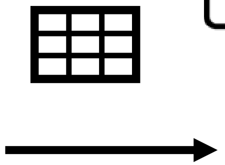

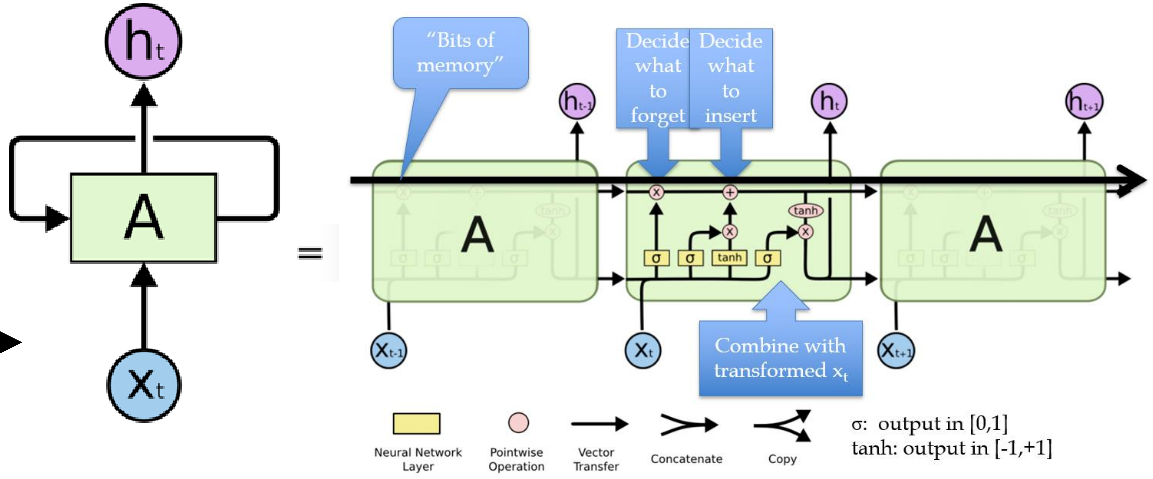

Deep Learning Workflow

Create and Access Datasets → **Preprocess and Transform Data** → **Develop Predictive Models** → **Accelerate and Deploy**

Create and Access Datasets

Data sources

Simulation and augmentation

Data Labeling

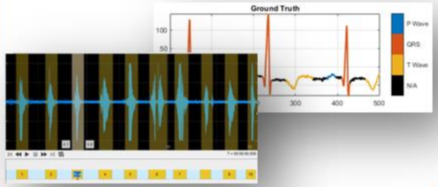

Preprocess and Transform Data

Pre-Processing

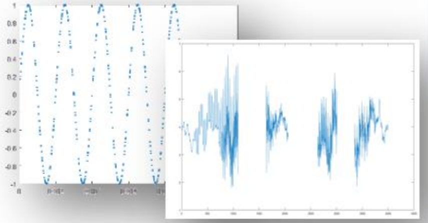

Transformation

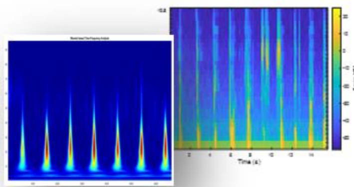

Feature extraction

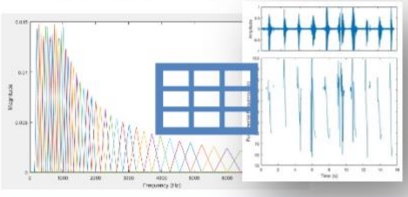

Develop Predictive Models

Import Reference Models/ Design from scratch

Hardware-Accelerated Training

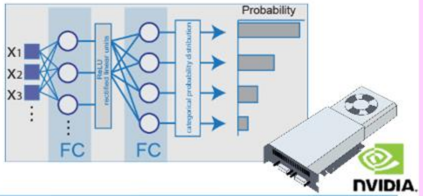

Analyze and tune hyperparameters

Accelerate and Deploy

Desktop Apps

Enterprise Scale Systems

Embedded Devices and Hardware

Deep Learning Workflow Challenges – Signals and Time Series

Data Labeling

Limited Data Availabiliy

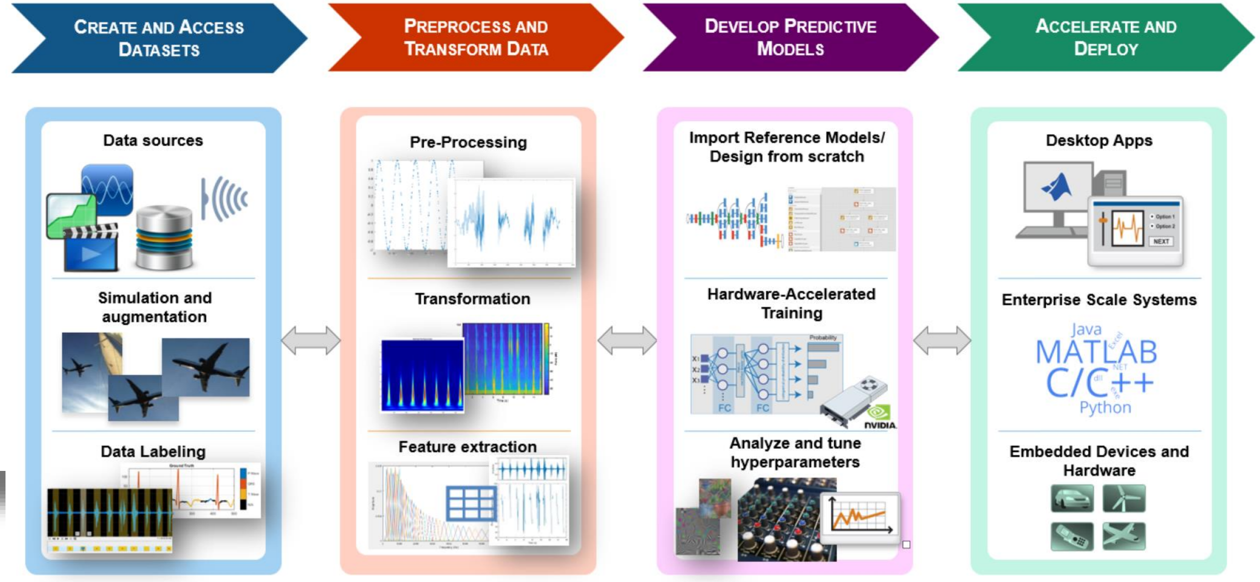

Deep learning models only as good as training data

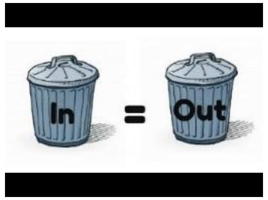

Domain-specific data processing desirable

Limited reference research

Deployment and Scaling to various platforms

Agenda

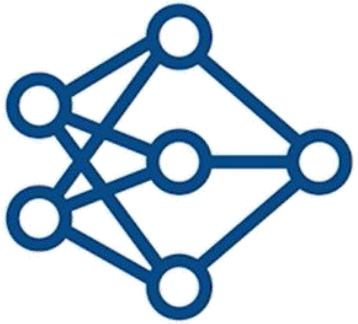

- Deep Learning – Basic ideas

- **Deep Learning Model Development for Signals, Time Series, and Text**
  - Data
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

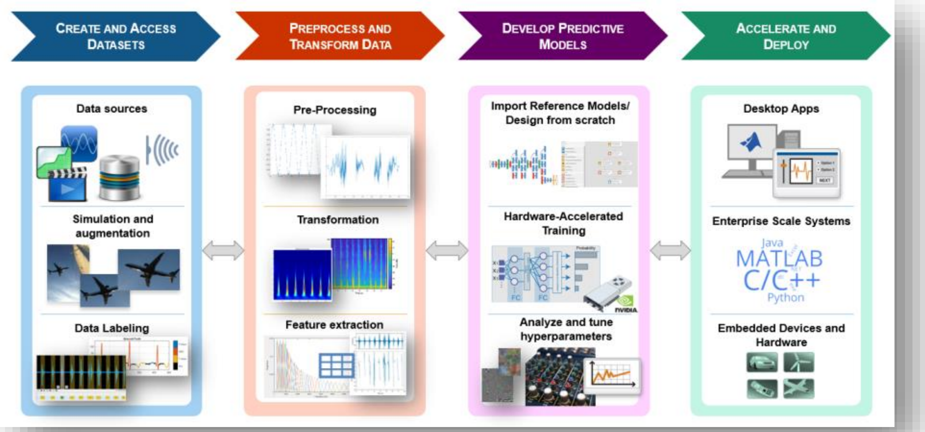

- Conclusions

Agenda

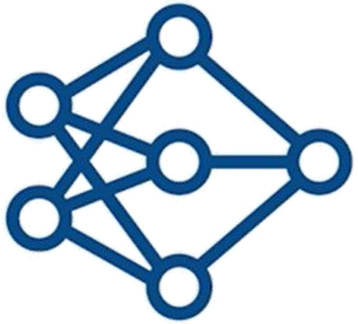

- Deep Learning – Basic ideas
- Deep Learning Model Development for Signals, Time Series, and Text
  - **Data**
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

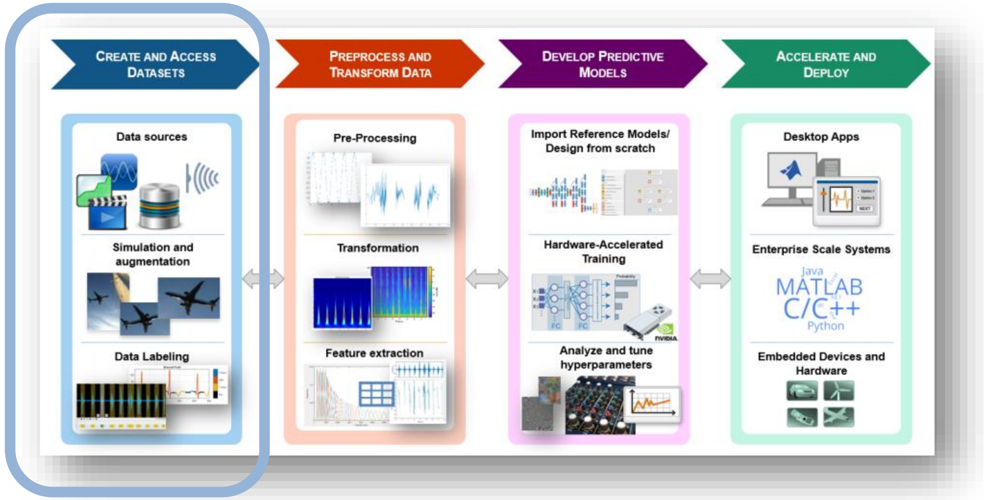

- Conclusions

Current Investments – Models vs. Data

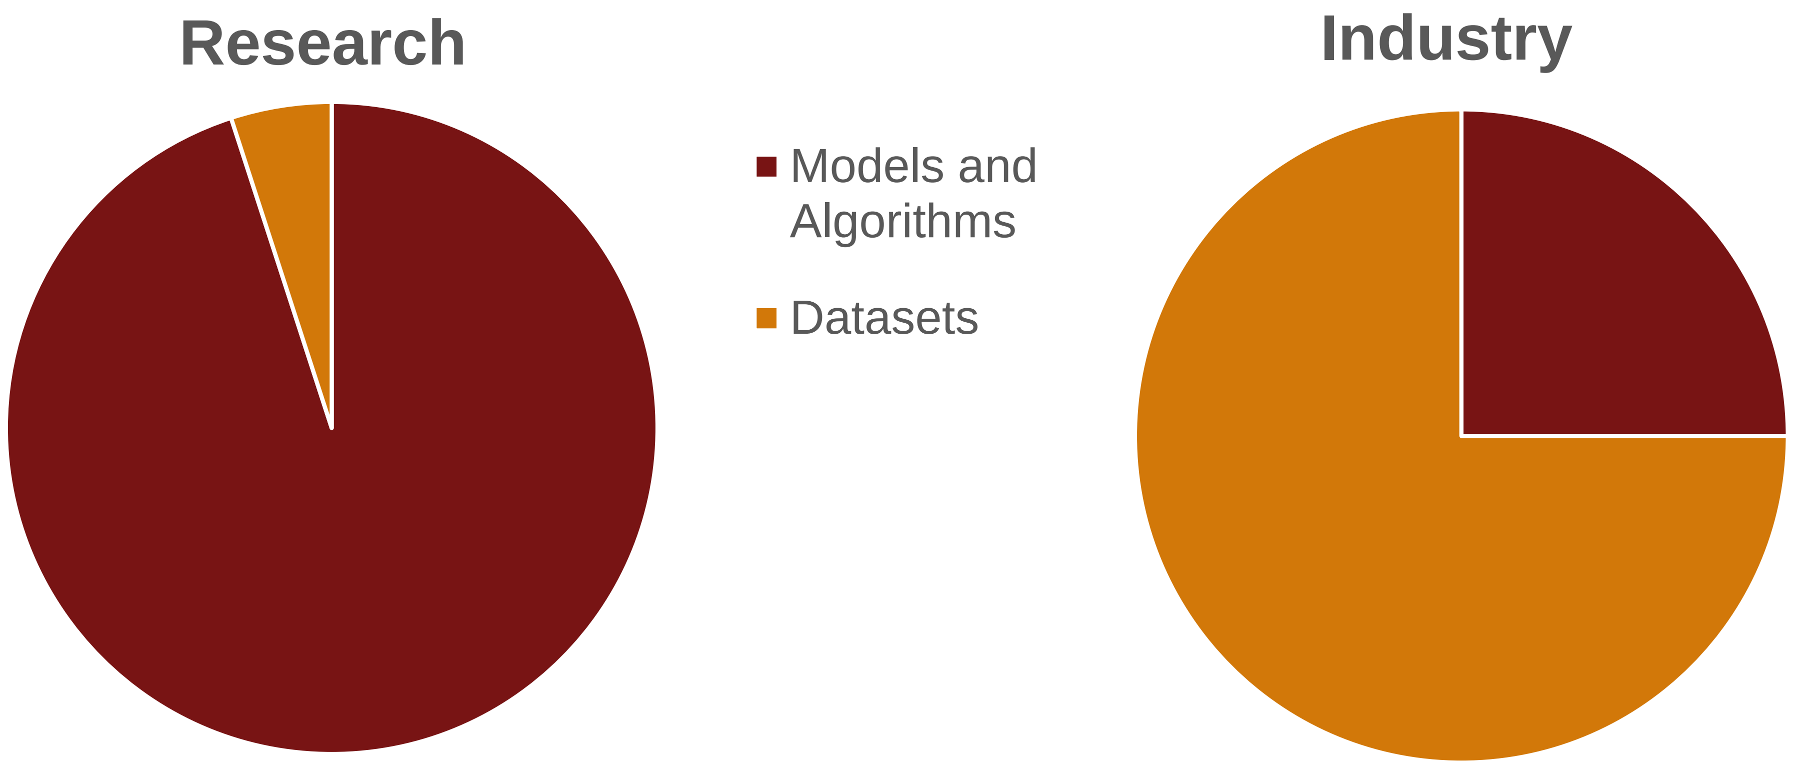

From "Troubleshooting deep neural networks" (Josh Tobin et al., Jan 2019)

What does a large dataset look

How to navigate, index, read (al

audioDatastore

imageDatastore

fileDatastore

Custom Datastores

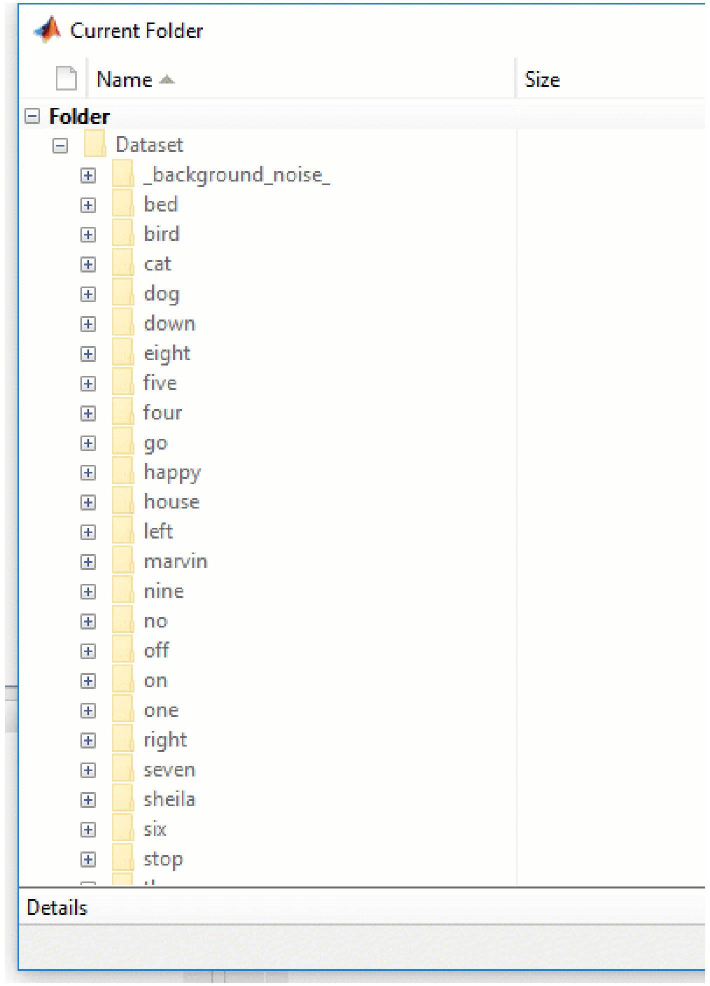

How to...

- Build a list of all data and labels?
- Review basic statistics about available data?
- Select data subsets without nested `for` loops, `dir`, `ls`, `what`, ... aplenty?
- Jointly read data and labels?
- Automatically distribute computations?

Label quality impacts model performance as much as the quality and quantity of the actual recordings

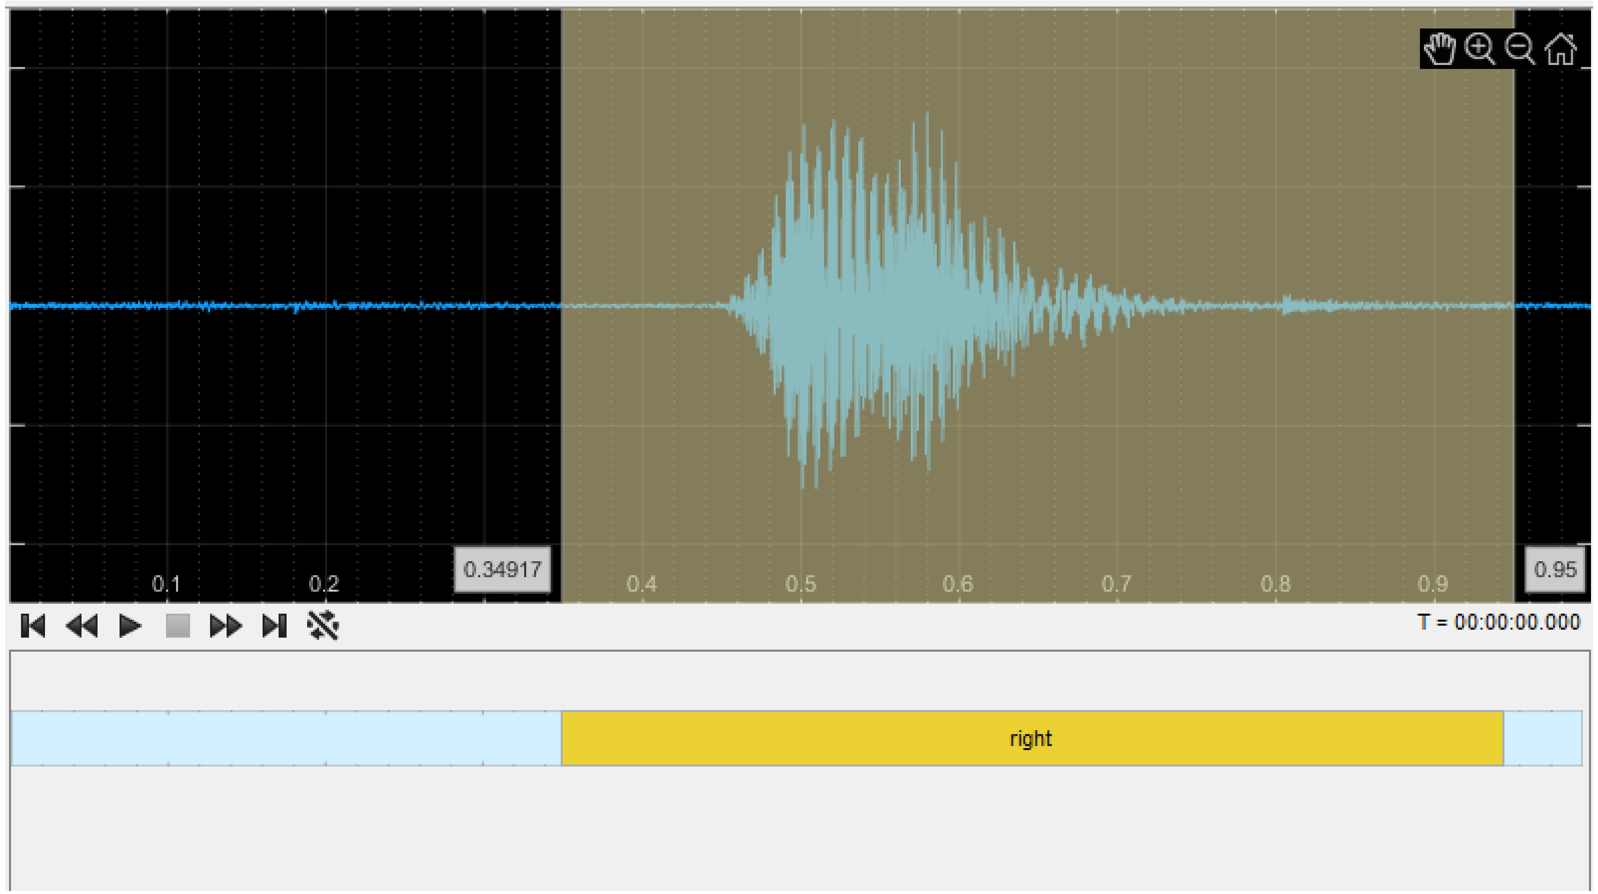

Use appropriate tools to help you label signals

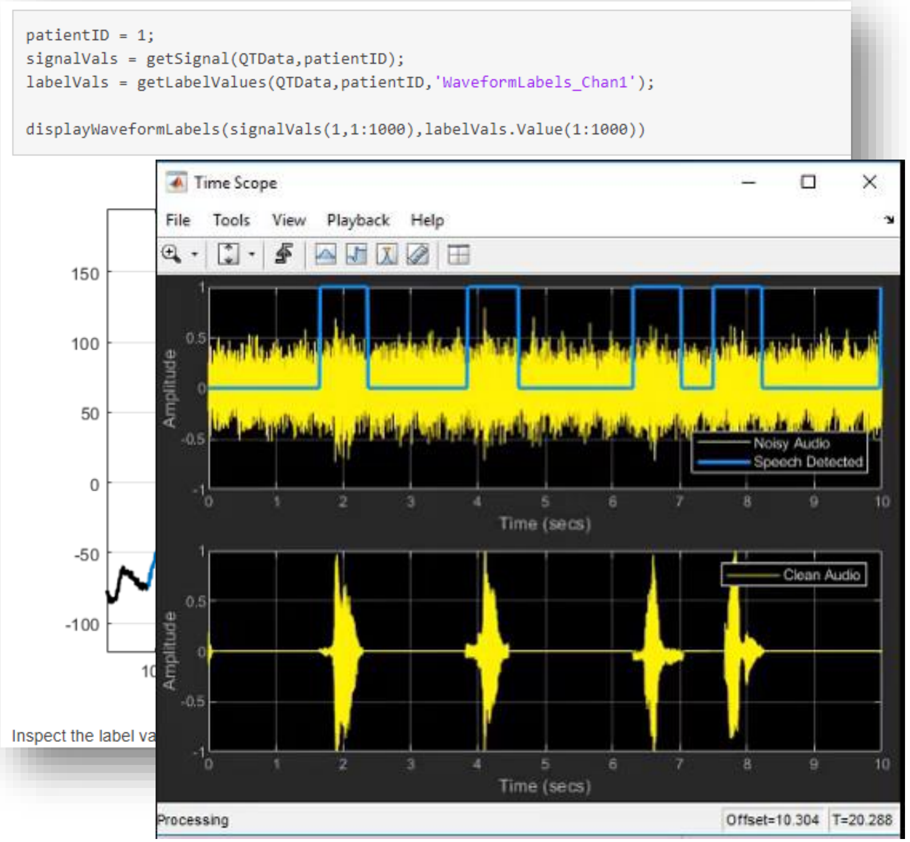

- Programmatically…

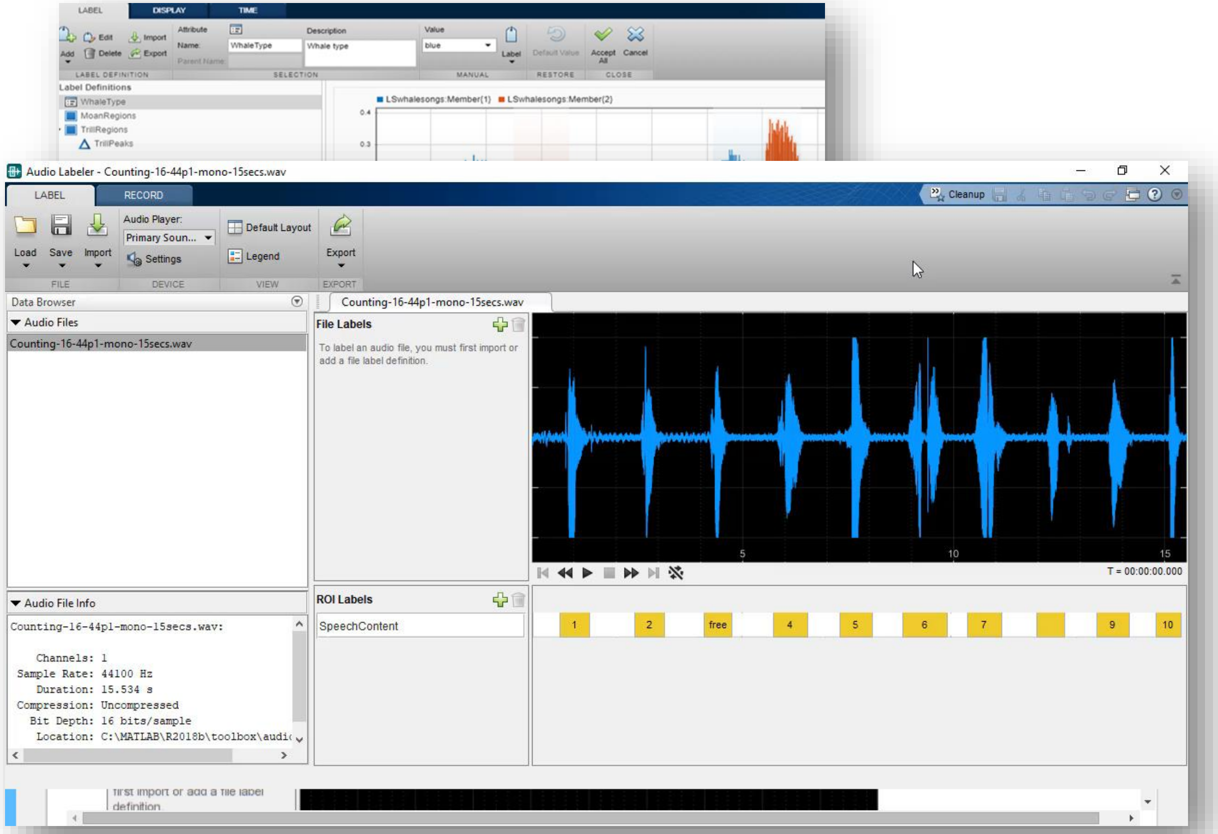

- … or via Apps

What if available data isn't enough?

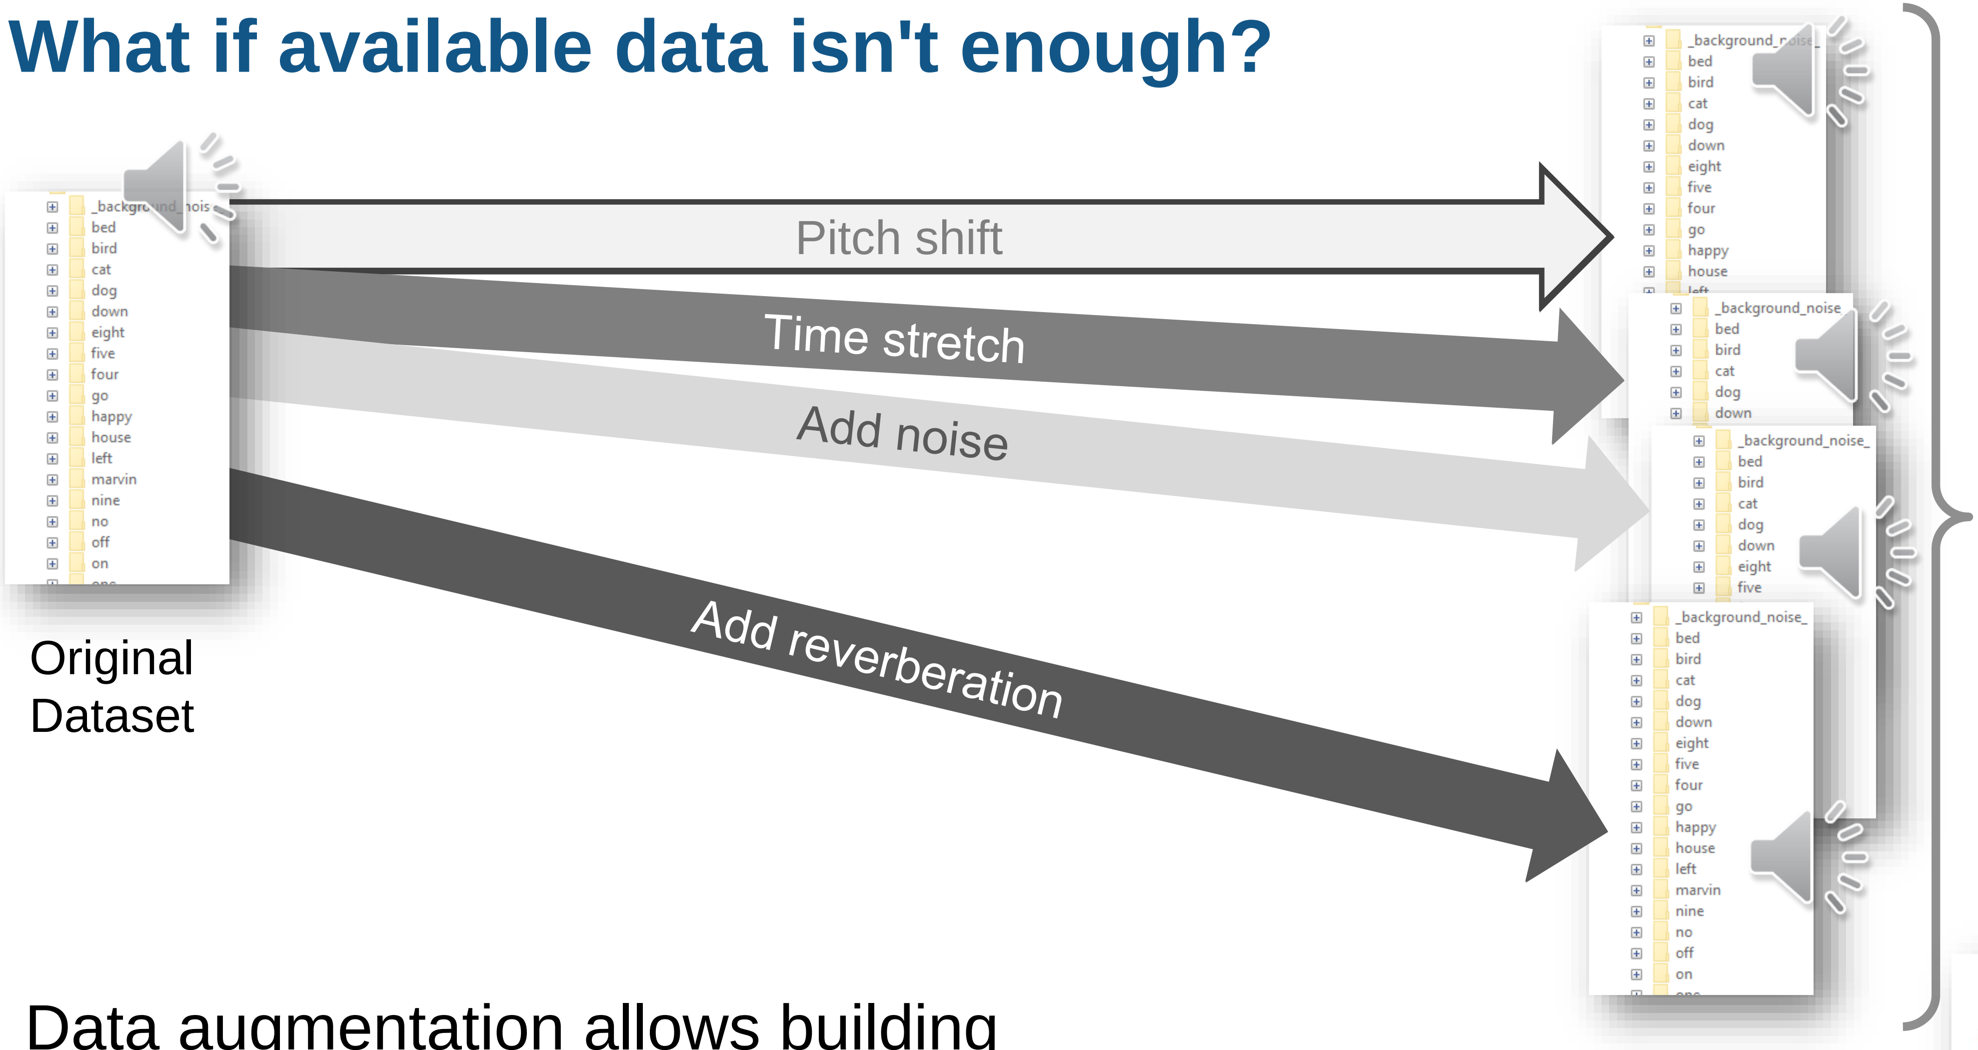

Augmented
Dataset

N times as
much data

Data augmentation allows building
more complex and more robust models

Simulation is key if recording and labelling real-world data is impractical or unreasonable – Communications Signals

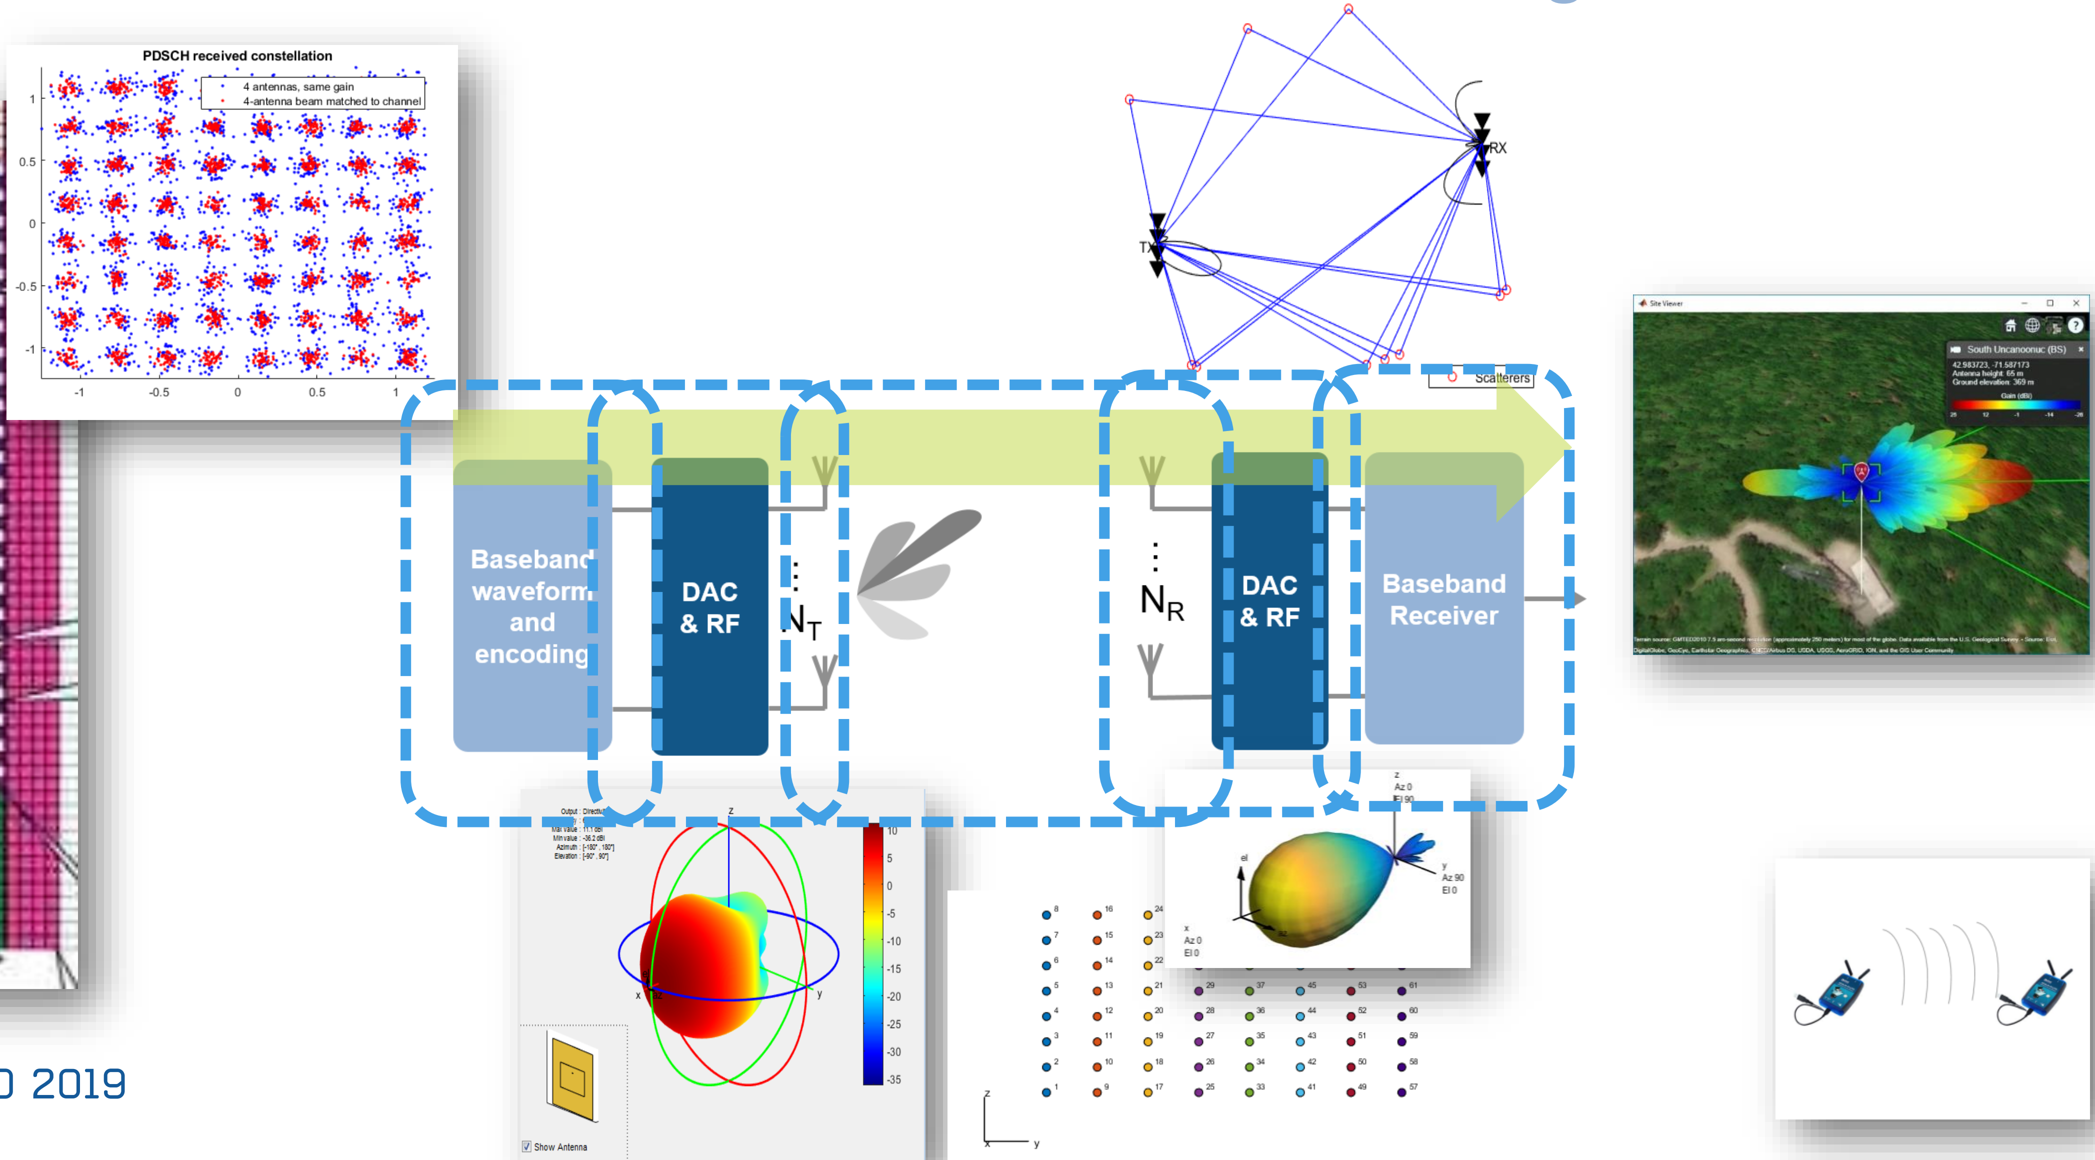

Simulation is key if recording and labelling real-world data is impractical or unreasonable – Radar Signals

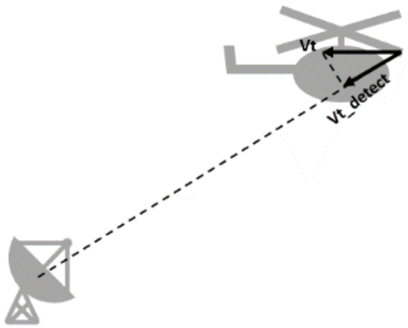

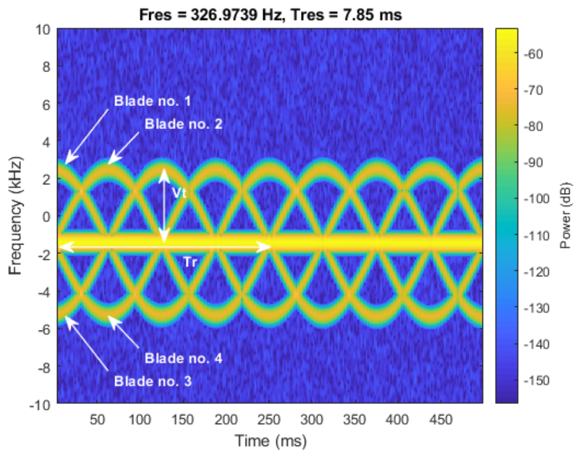

Radar Target Simulation

Micro-Doppler Analysis

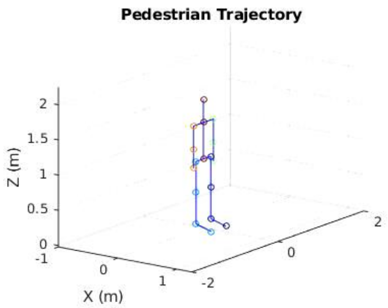

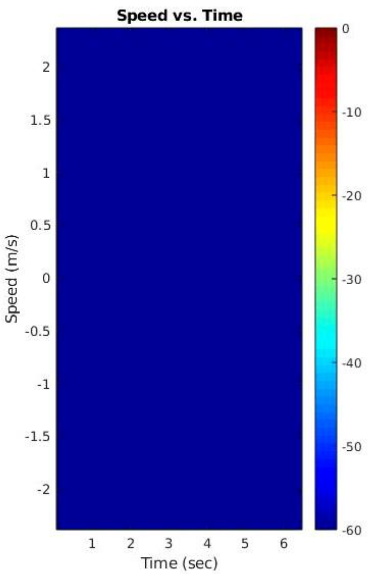

Agenda

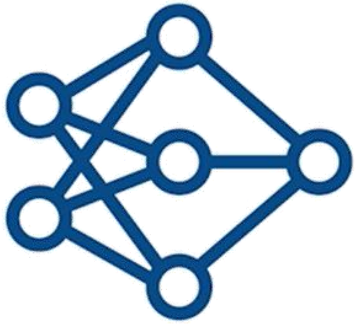

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time Series, and Text
  - Data
  - **Processing and transformation**
  - Model design and optimization
  - Acceleration, prototyping, and deployment

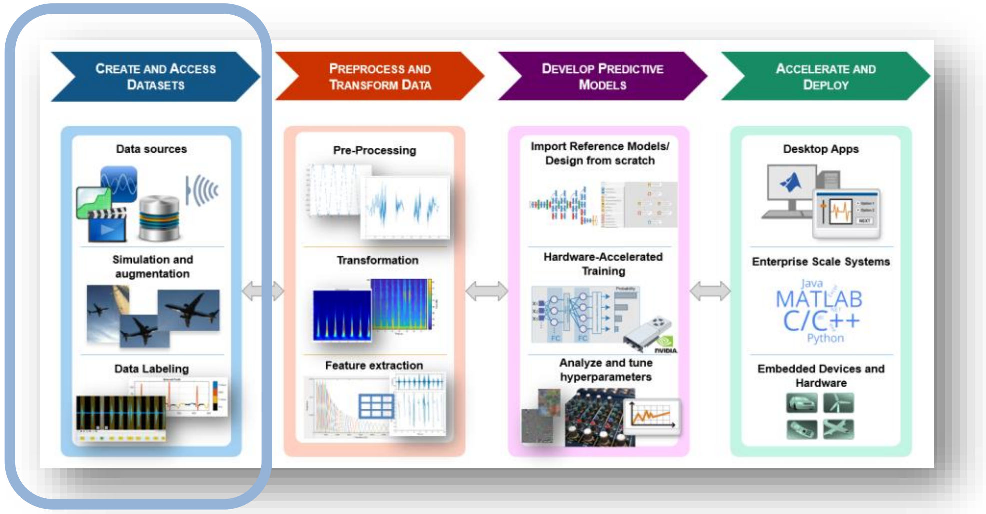

- Conclusions

Common types of network architectures used in signal processing and text analytics applications

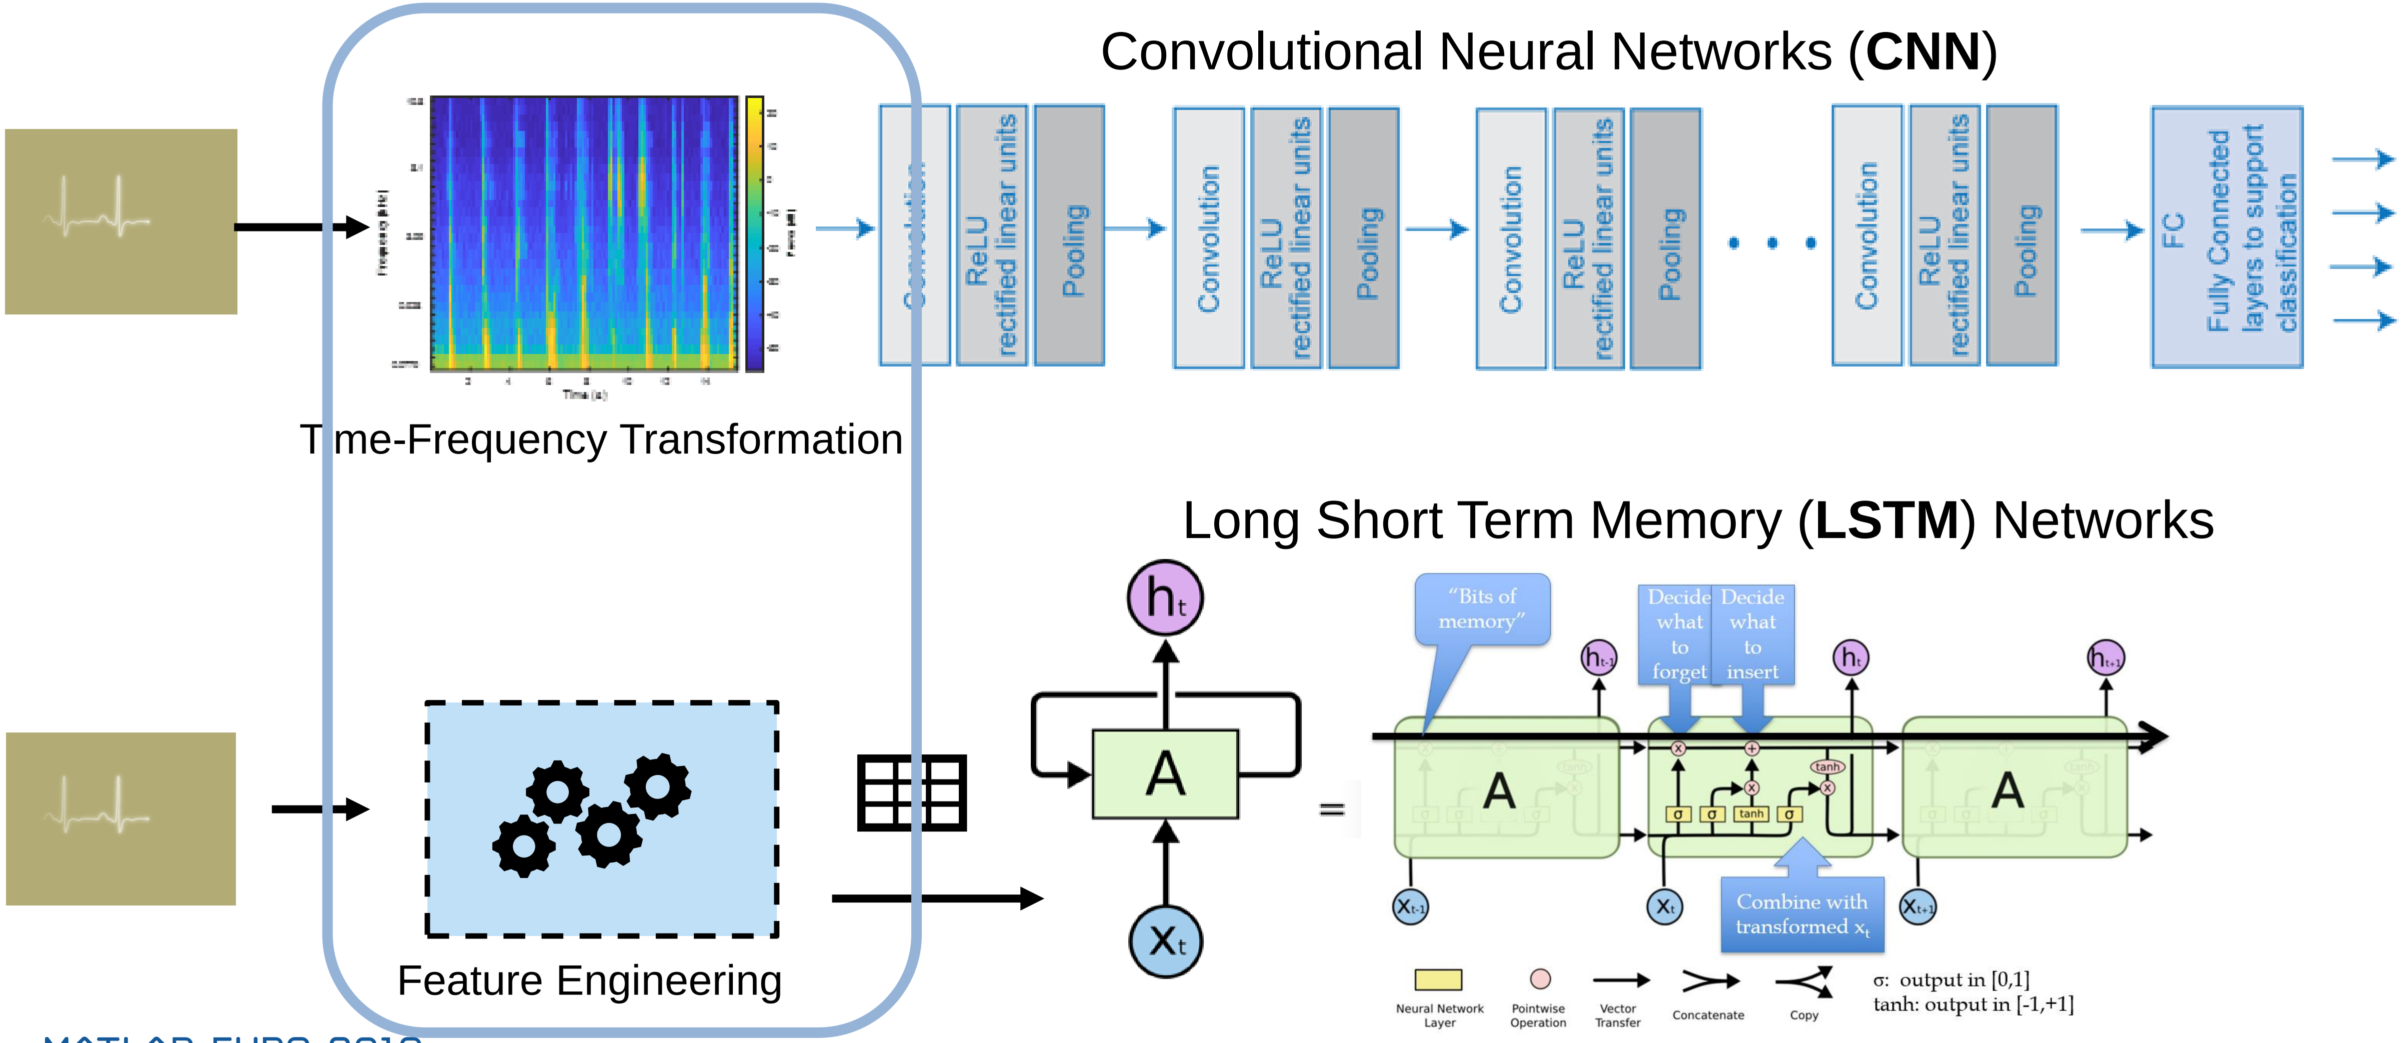

Time-Frequency Transformations

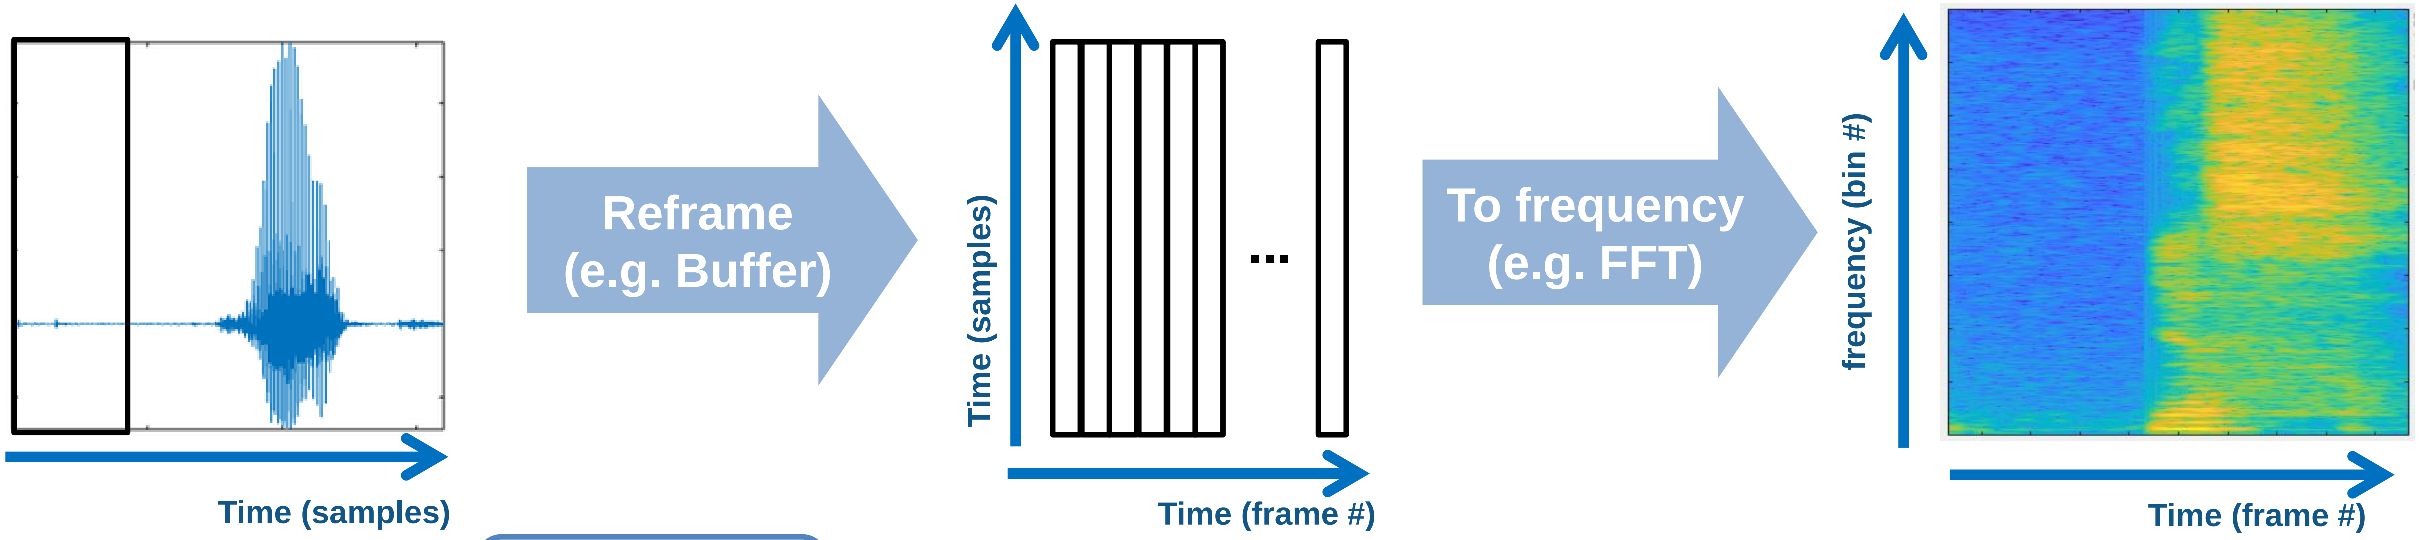

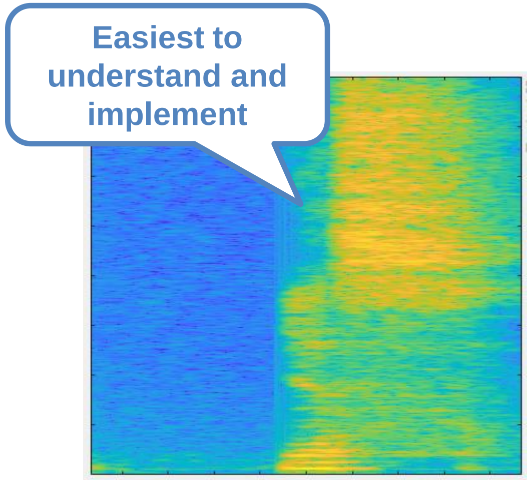

Basic spectrogram

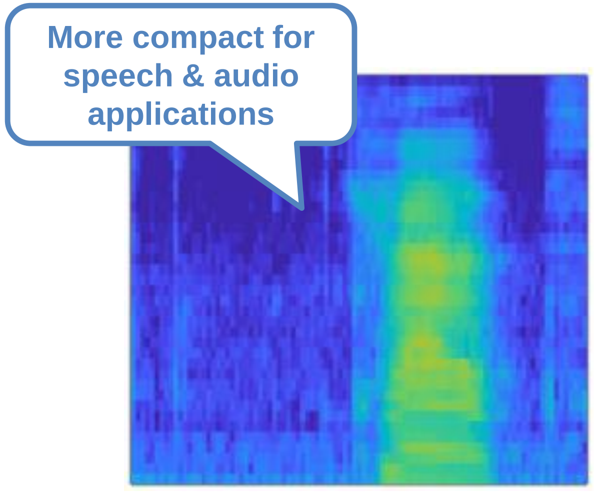

Perceptually-spaced (e.g. Mel, Bark) Spectrogram

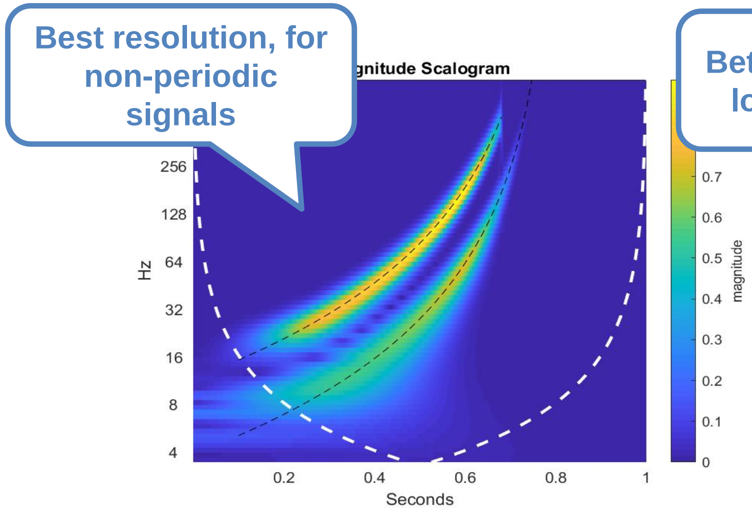

Wavelet scalogram

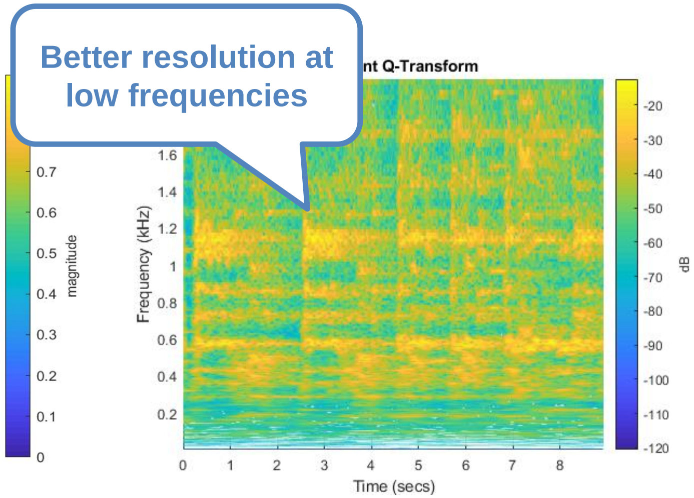

Constant Q transform

Extracting Features from Signals: Application-Agnostic Examples

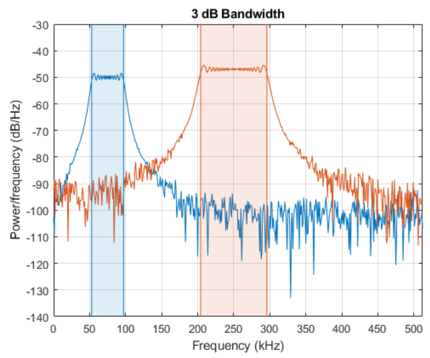

BW measurements

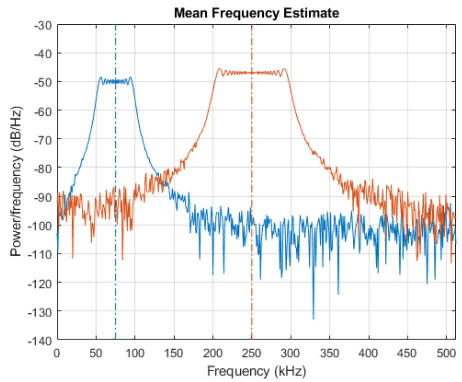

Spectral statistics

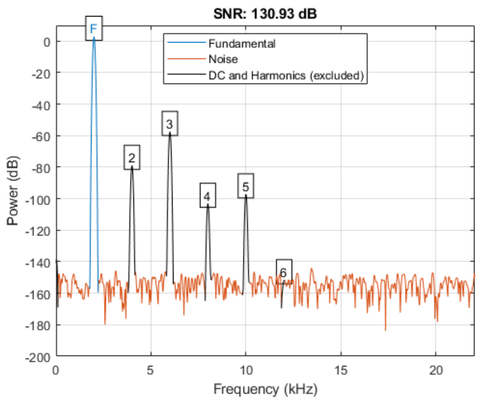

Harmonic analysis

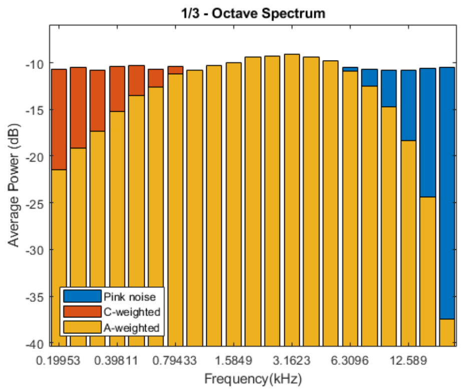

Octave spectrum

Frequency domain

Time domain

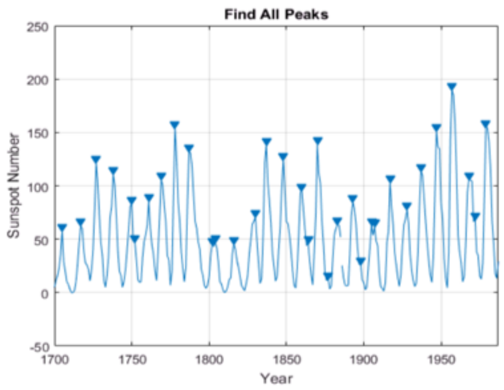

Find peaks

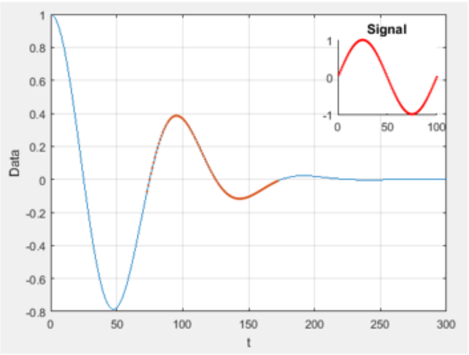

Find signal patterns

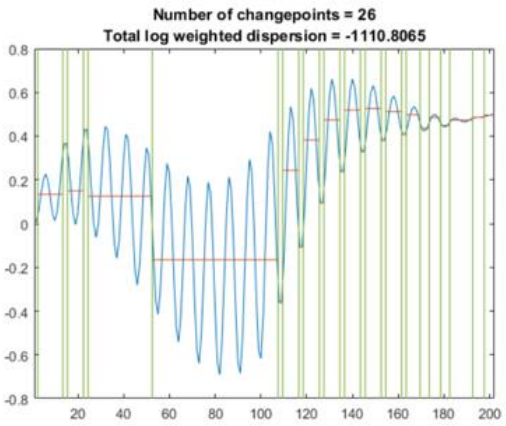

Detect change points

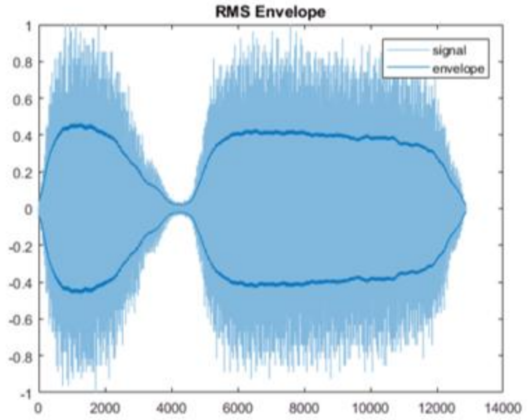

Find signal envelope

Domain-Specific Features and Transformations – Examples

Speech and Audio

- MFCC
- GTCC
- MDCT
- Pitch, harmonicity
- Spectral shape descriptors
- ...

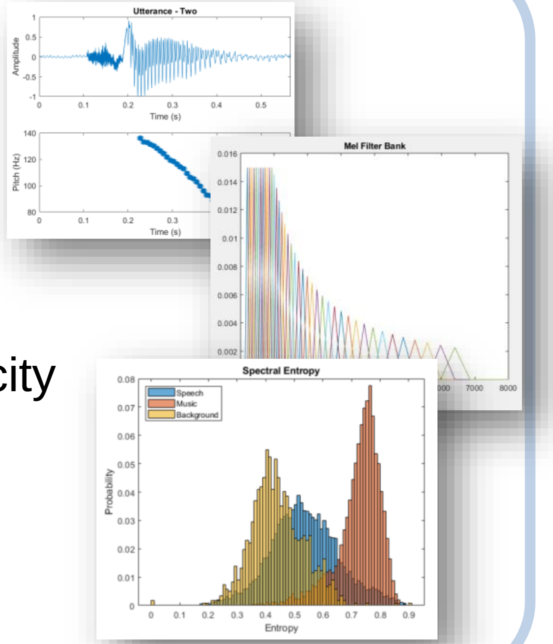

Navigation and Sensor Fusion

- Orientation
- Height
- Position

from

- Acceleration, angular velocity
- Magnetic field
- GPS reading

- Multi-object tracking
- …

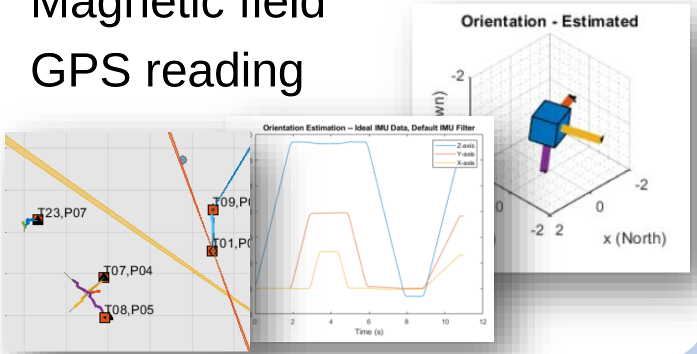

Text Analytics

- Train Word Embeddings
- Word2Vec
- Topic Modeling
- ...

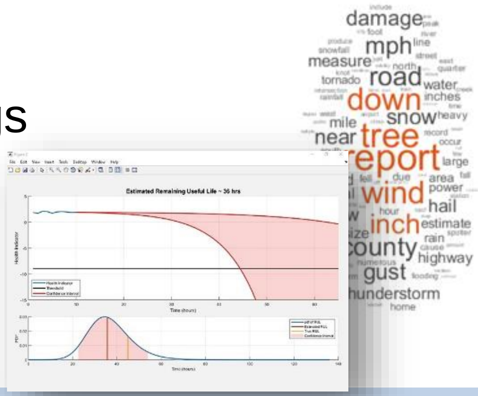

Radar

- Micro-Doppler analysis
- Range-Doppler processing
- Synthetic aperture imaging
- Spectral analysis
- Waveform ambiguity
- ...

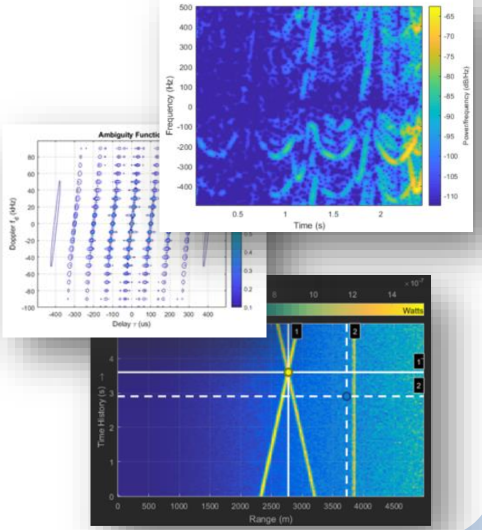

Automated Feature Extraction: Wavelet Scattering

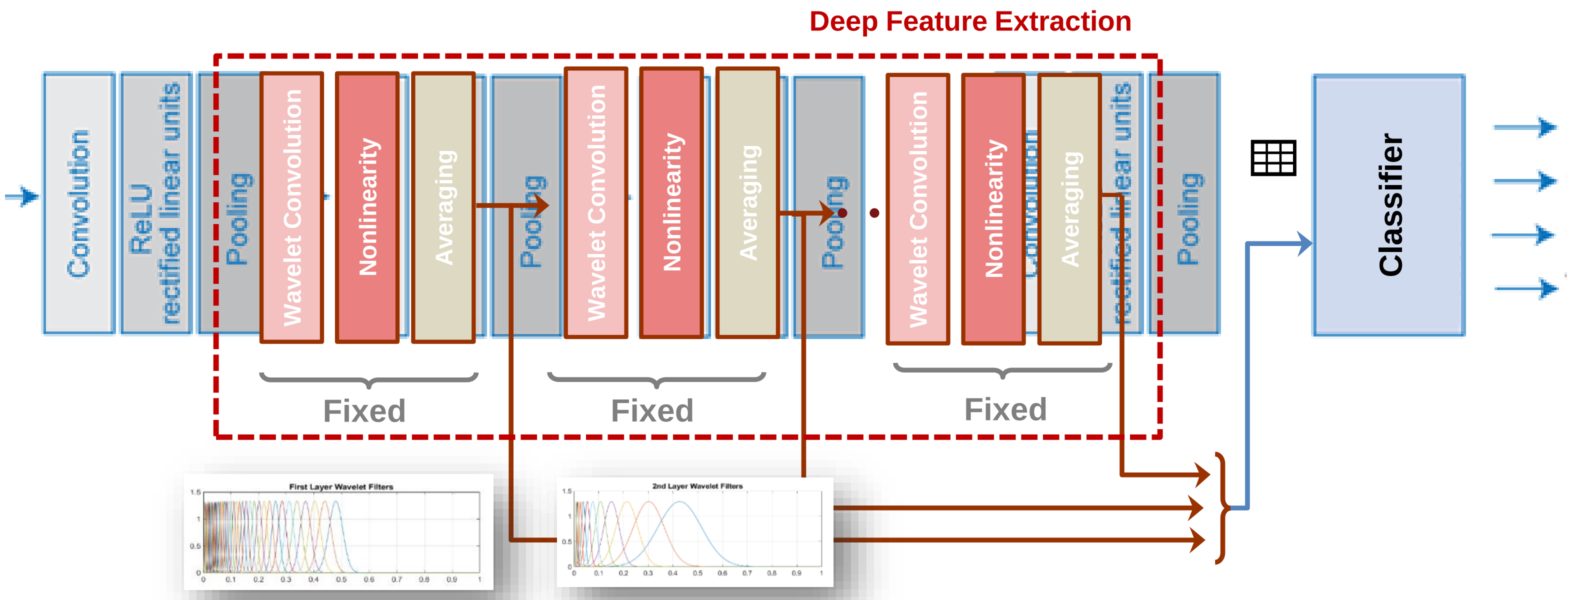

- Can relieve requirements on amount of data and model complexity
  - Featured in leader-boards a number of research competitions

- Framework for extracting features [1]

[1] Joan Bruna, and Stephane Mallat, P. 2013. Invariant Scattering Convolution Networks. IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 35, No. 8, pp. 1872-1886.

Agenda

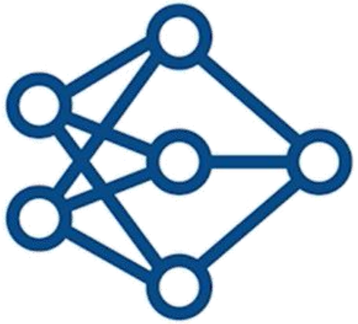

- Deep Learning – Basic ideas
- Deep Learning Model Development for Signals, Time Series, and Text
  - Data
  - Processing and transformation
  - **Model design and optimization**
  - Acceleration, prototyping, and deployment

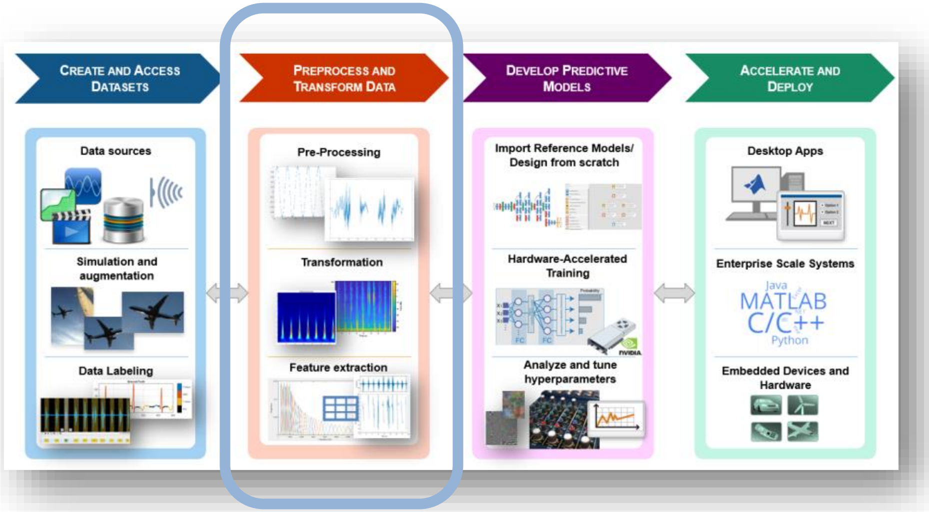

- Conclusions

Developing Deep Learning Models

Design Network

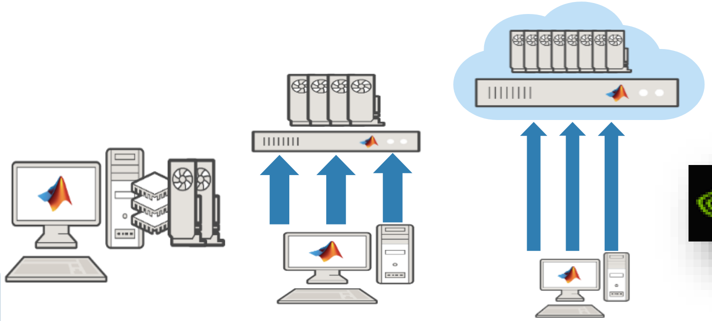

Accelerate Training

NVIDIA GPU CLOUD

MATLAB as a container on NGC

Train

Design

Optimize

Model Exchange

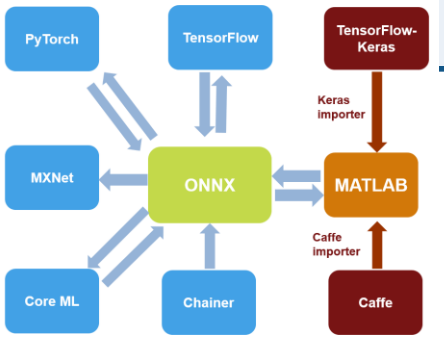

Pre-trained Networks

AlexNet PRETRAINED MODEL
VGG-16 PRETRAINED MODEL
ResNet-50 PRETRAINED MODEL
ONNX Converter MODEL CONVERTER
Caffe IMPORTER
GoogLeNet PRETRAINED MODEL
TensorFlow-Keras IMPORTER
Inception-v3 MODELS

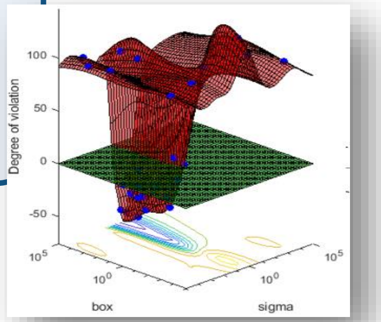

Bayesian Hyperparameter Optimization

Exchange Models With Deep Learning Frameworks

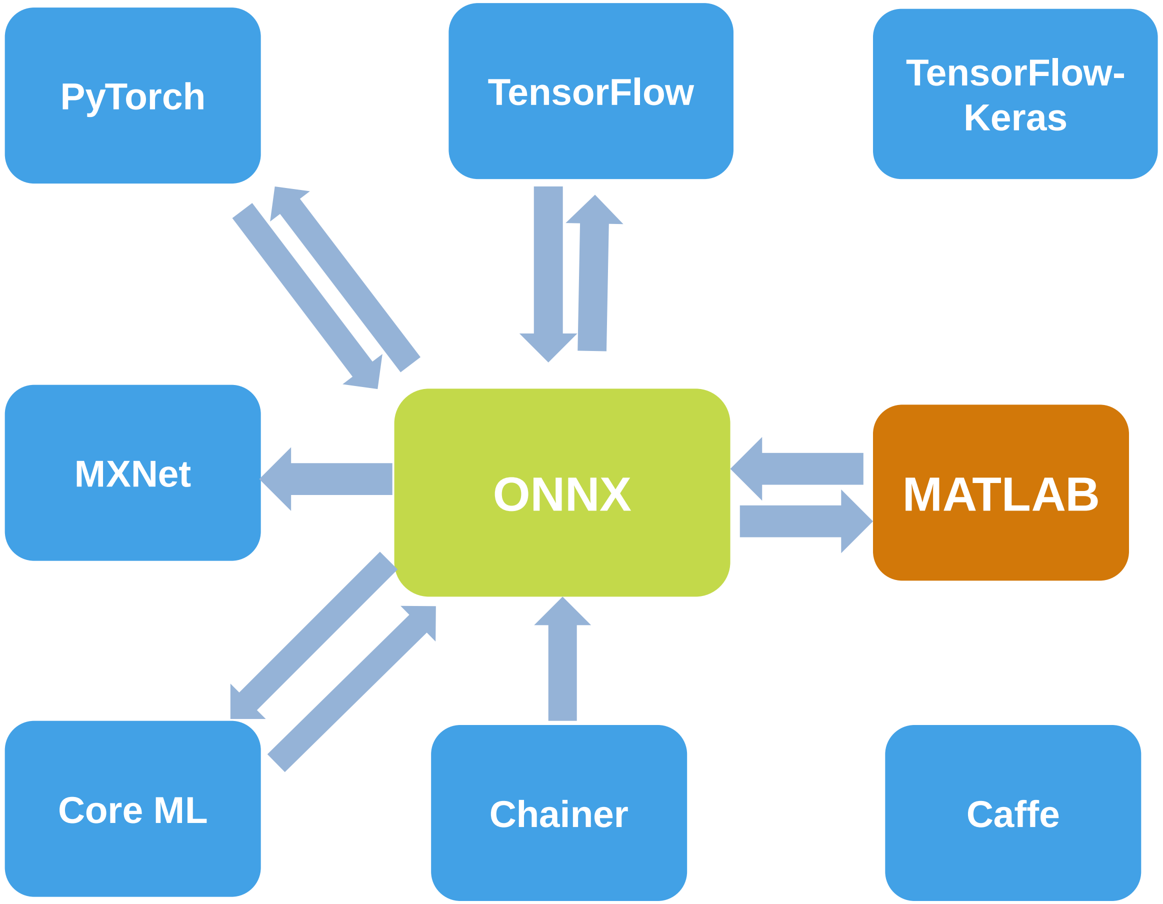

ONNX = Open Neural Network Exchange Format

Exchange Models With Deep Learning Frameworks

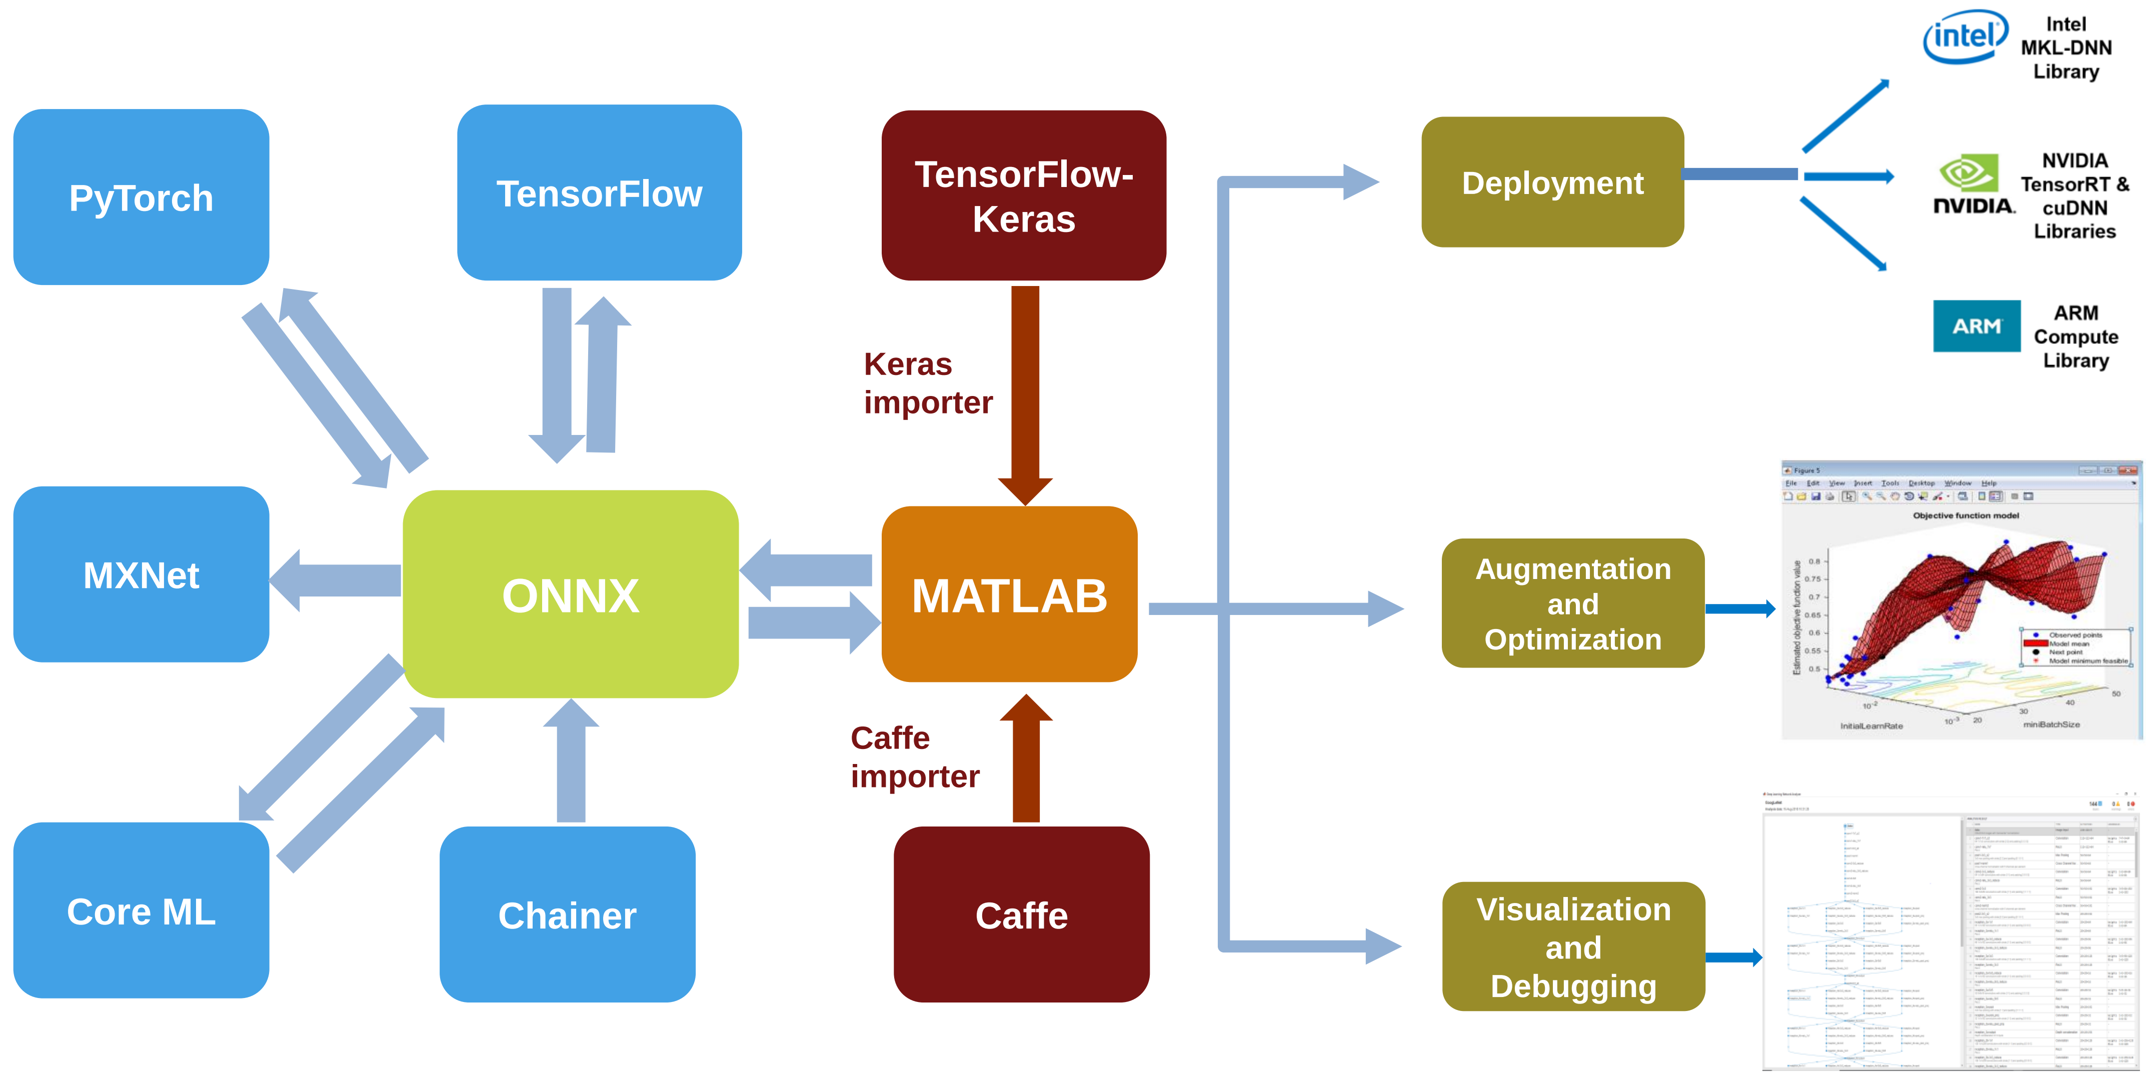

ONNX = Open Neural Network Exchange Format

Agenda

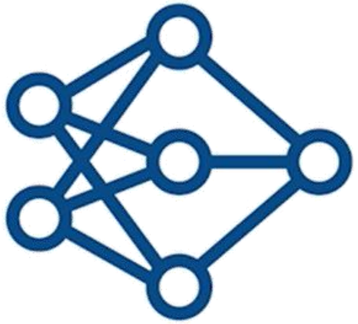

- Deep Learning – Basic ideas
- Deep Learning Model Development for Signals, Time Series, and Text
  - Data
  - Processing and transformation
  - Model design and optimization
  - **Acceleration, prototyping, and deployment**

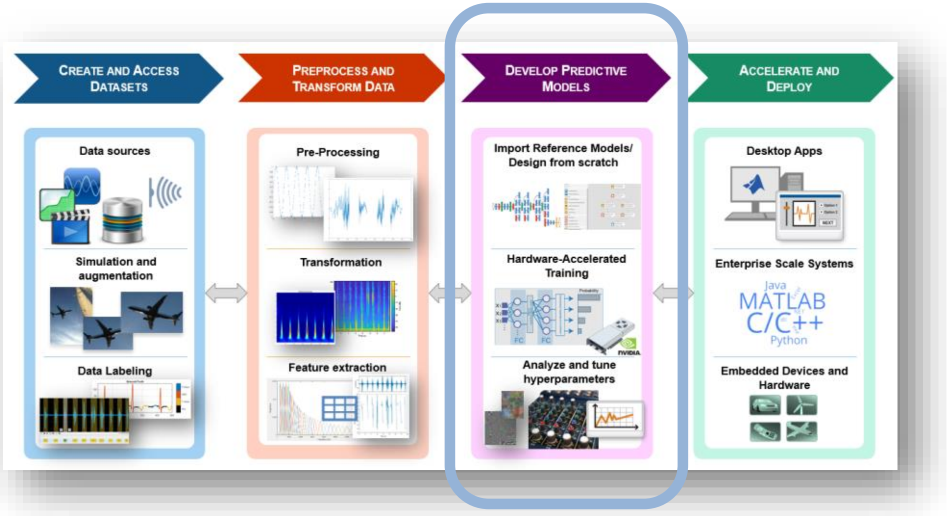

- Conclusions

Deployment and Scaling for A.I.

MATLAB

Embedded Devices

Enterprise Systems

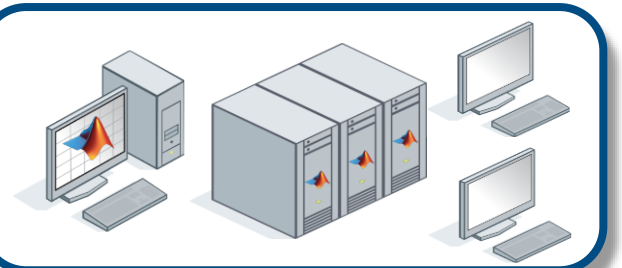

Deploying Deep Learning Models for Inference

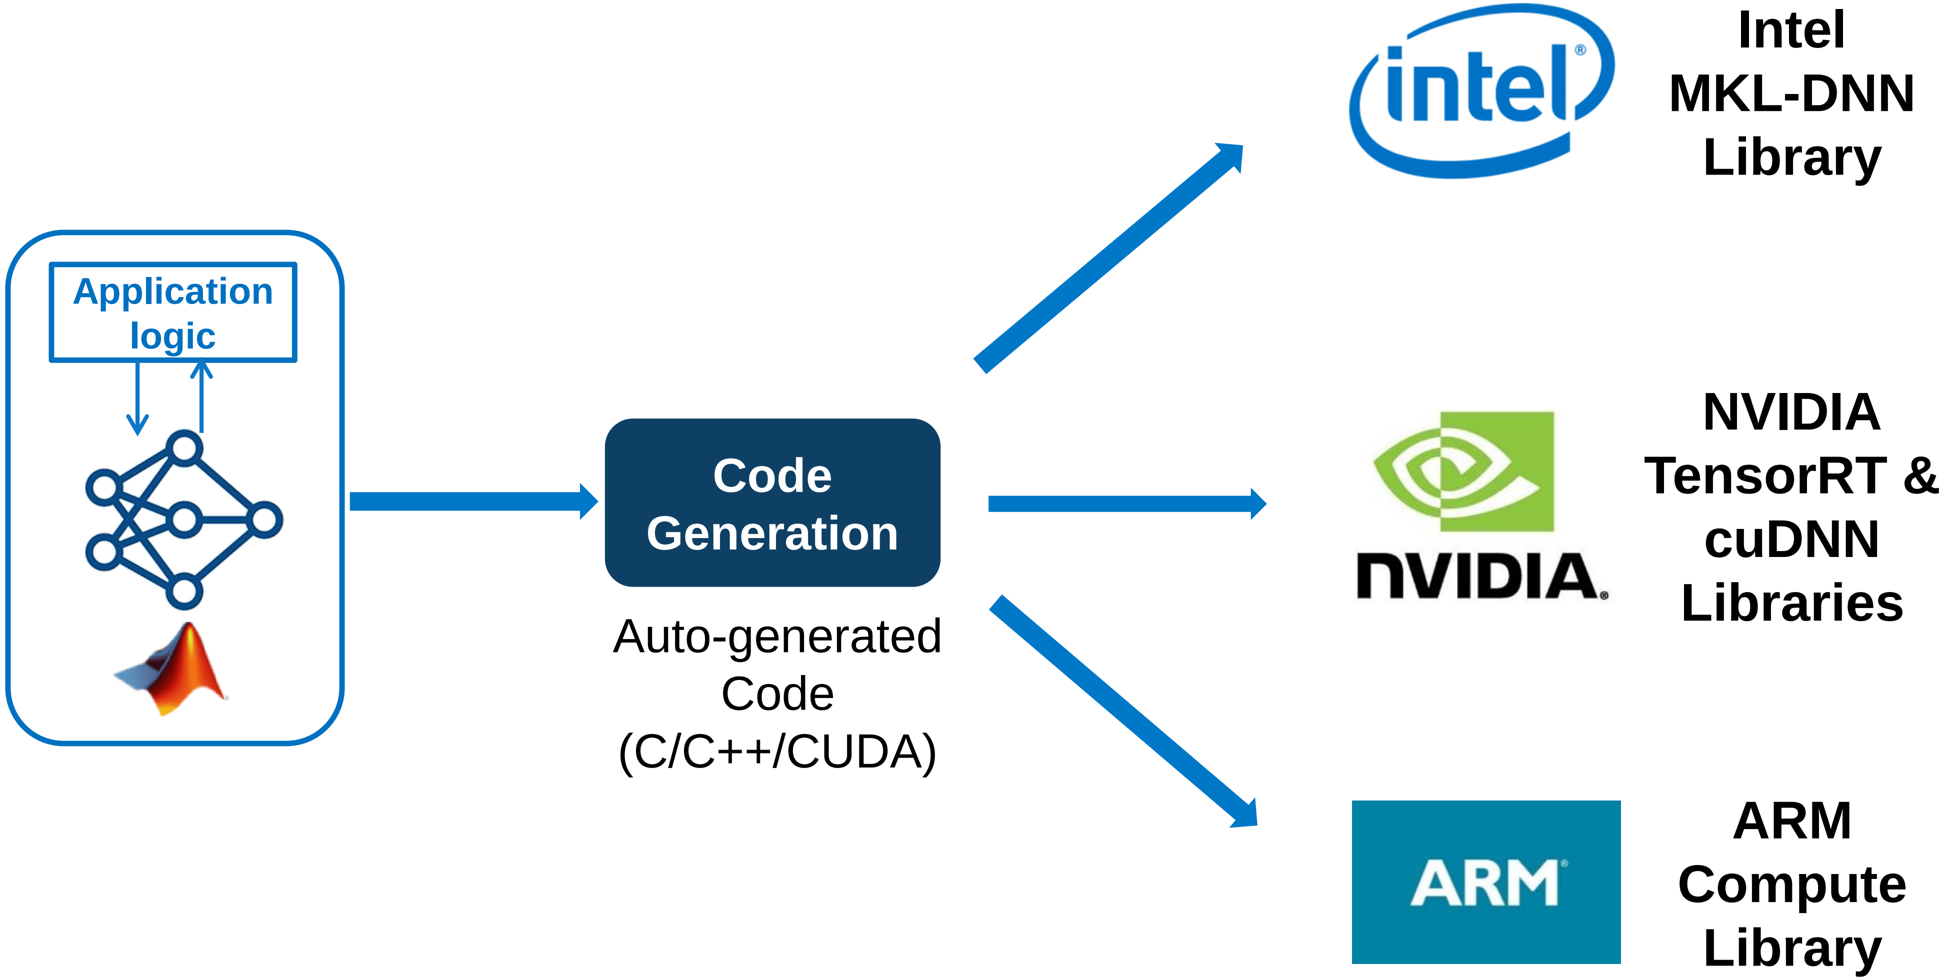

Enterprise Deployment

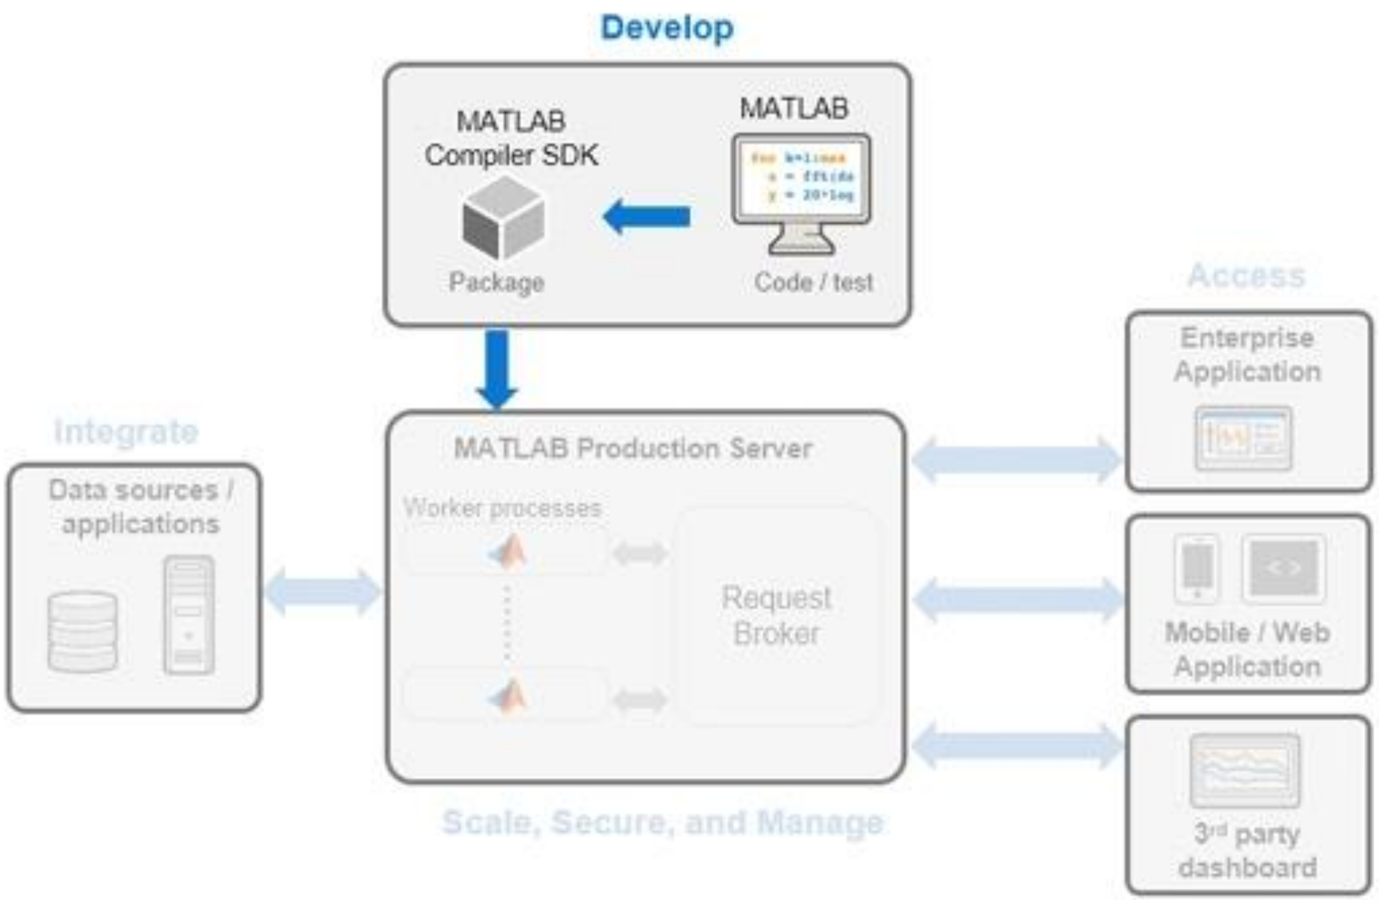

Deployment to the cloud with MATLAB Compiler and MATLAB Production Server

Agenda

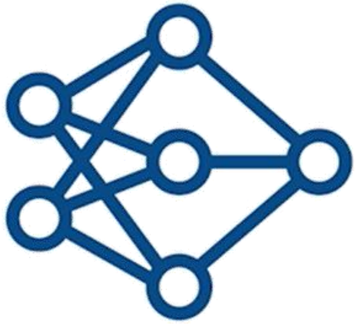

- Deep Learning – Basic ideas

- Deep Learning Model Development for Signals, Time Series, and Text
  - Data
  - Processing and transformation
  - Model design and optimization
  - Acceleration, prototyping, and deployment

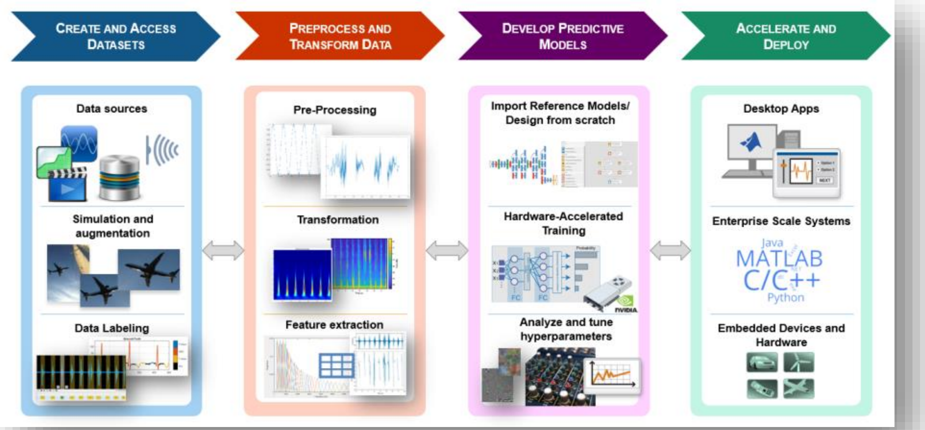

- **Conclusions**

Deep Learning Workflow Challenges – Signals and Time Series

Data-labeling Apps and Examples

Augmentation and simulation algorithms

Deployment and Scaling to various platforms

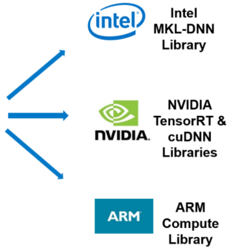

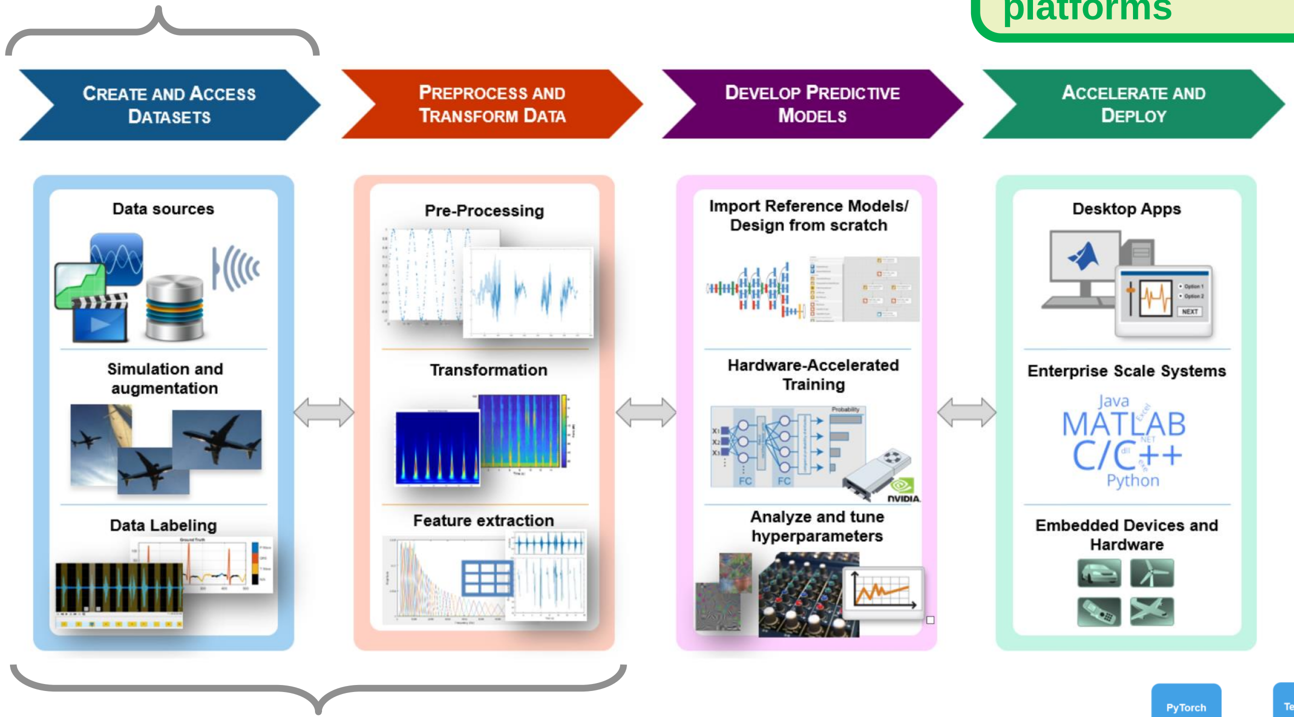

Deep learning models only as good as training data

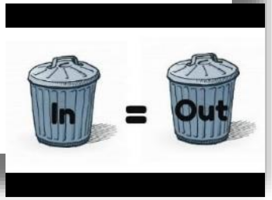

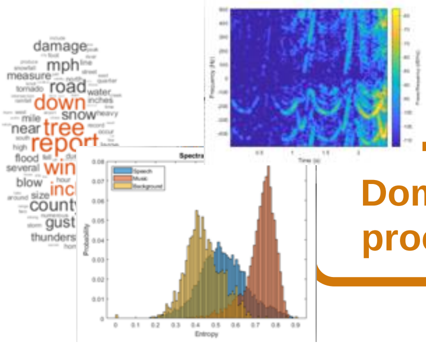

Dom
proc

Application-specific algorithms and tools

Collaboration in the AI ecosystem

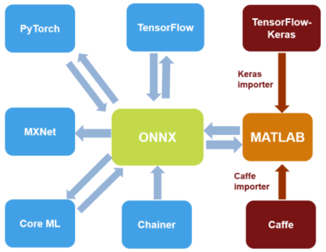

Domain-Specific Features and Transformations – Examples

Audio

Speech Command Recognition
Voice Activity Detection in Noise
Denoise Speech
Classify Gender

Time-Series and Text

Classify Time Series Using Wavelet Analysis
Sequence-to-Sequence Classification
Classify Text Data Using LSTMs
Classify Text Data Using CNNs

Signal

Music Genre Classification
Human Activity Recognition
ECG Signal Classification
Waveform Segmentation

Comms and Radar

Radar Waveform Classification
Modulation Classification

What next?

- Deep Learning Onramp
- Other talks:
  - Pixels to Features to Models
  - Automated Driving System Design
- Demo stands:
  - Deep Learning and Reinforcement Learning
  - Driverless – Science Museum exhibition stand
- Application Engineer support

Back up

Summary - Deep learning workflow in MATLAB